

Continuing Commentary

- neuroendocrine response patterns. *Psychoneuroendocrinology* 25:577–91. [JvH]
- van Honk, J., Tuiten, A., Verbaten, R., van den Hout, M., Koppeschaar, H., Thijsen, J. & de Haan, E. (1999) Correlations among salivary testosterone, mood, and selective attention to threat in humans. *Hormones and Behavior* 36:17–24. [JvH]
- Wood, R. I. (1996) Functions of the steroid-responsive neural network in the control of male hamster sexual behavior. *Trends in Endocrinology and Metabolism* 7:338–44. [JvH]

Allan Mazur & Alan Booth have declined to respond to the above continuing commentaries.

Commentary on Lawrence W. Barsalou (1999). Perceptual symbol systems. *BBS* 22(4):577–660.

Abstract of the original article: Prior to the twentieth century, theories of knowledge were inherently perceptual. Since then, developments in logic, statistics, and programming languages have inspired amodal theories that rest on principles fundamentally different from those underlying perception. In addition, perceptual approaches have become widely viewed as untenable because they are assumed to implement recording systems, not conceptual systems. A perceptual theory of knowledge is developed here in the context of current cognitive science and neuroscience. During perceptual experience, association areas in the brain capture bottom-up patterns of activation in sensory-motor areas. Later, in a top-down manner, association areas partially reactivate sensory-motor areas to implement perceptual symbols. The storage and reactivation of perceptual symbols operates at the level of perceptual components – not at the level of holistic perceptual experiences. Through the use of selective attention, schematic representations of perceptual components are extracted from experience and stored in memory (e.g., individual memories of *green*, *purr*, *hot*). As memories of the same component become organized around a common frame, they implement a simulator that produces limitless simulations of the component (e.g., simulations of *purr*). Not only do such simulators develop for aspects of sensory experience, they also develop for aspects of proprioception (e.g., *lift*, *run*) and introspection (e.g., *compare*, *memory*, *happy*, *hungry*). Once established, these simulators implement a basic conceptual system that represents types, supports categorization, and produces categorical inferences. These simulators further support productivity, propositions, and abstract concepts, thereby implementing a fully functional conceptual system. Productivity results from integrating simulators combinatorially and recursively to produce complex simulations. Propositions result from binding simulators to perceived individuals to represent type-token relations. Abstract concepts are grounded in complex simulations of combined physical and introspective events. Thus, a perceptual theory of knowledge can implement a fully functional conceptual system while avoiding problems associated with amodal symbol systems. Implications for cognition, neuroscience, evolution, development, and artificial intelligence are explored.

Amodal or perceptual symbol systems: A false dichotomy?

W. Martin Davies

Faculty of Economics and Commerce, University of Melbourne, Victoria 3010, Australia. wmdavies@unimelb.edu.au


Abstract: Although Barsalou is right in identifying the importance of perceptual symbols as a means of carrying certain kinds of content, he is wrong in playing down the inferential resources available to amodal symbols. I argue that the case for perceptual symbol systems amounts to a false dichotomy and that it is feasible to help oneself to both kinds of content as extreme ends on a content continuum. The continuum thesis I advance argues for the inferential content at one end and perceptual content at the other. In between the extremes, symbols might have aspects that are either perceptual or propositional-linguistic in character. I argue that this way of characterising the issue preserves the good sense of Barsalou's recognition of perceptual representations and yet avoids the tendency to minimise the gains won with symbolic representations vital to contemporary cognitive science.


In his target article, Lawrence Barsalou (1999t) has argued the case for a perceptual symbol systems approach in cognitive science on the grounds that the current orthodoxy, the amodal approach, has too many flaws. Barsalou identifies six central problems for amodalism: (1) there is no evidence that amodal symbols exist; (2) neuroscientific evidence points to activity in sensory motor regions of the brain on certain tasks; (3) amodal symbols have problems coping with representing certain cognitive processes such as spatio-temporal knowledge; (4) there is no satisfactory way in which amodal symbols can be mapped onto the perceptual states that caused them (the “transduction” problem); (5) there is no clear account of the manner in which amodal symbols can be mapped back onto perceptual states in the world (the “symbol grounding” problem); and finally, (6) amodal symbols are power-

fully explanatory and predictive in a post hoc fashion but not in any other way – a feature that makes them unfalsifiable.

Many of these difficulties can be levelled just as easily at the perceptual symbol approach, I suspect. Even some of the strongest evidence for perceptual imagery (e.g., Kosslyn 1994; Lang 1979; Shepard & Metzler 1971) suggest only principled support for the existence of imagery, not direct evidence. Equally, while it can also be fairly said that amodal symbols do not handle many aspects of cognition, so it is also true that perceptual symbols cannot handle other aspects, or do so with great difficulty. As for the claim about falsifiability (sect. 1.2.2 of the target article), in the current climate this seems equally true of perceptual symbols, and the debate so far is zero gain for either camp.

As for the neuroscientific evidence (sects. 2.1, 2.2, and 2.3 of Barsalou 1999t), it can hardly be argued that this is unambiguous evidence for either view. We surely know very little about the brain. Only if one conflates *correlations* and *causes* is there any hope of identifying certain brain processes with the mechanisms that are their supposed casual antecedents. Spring is correlated with the presence of bees in the air, but it would be a mistake to identify the two or to ground one in terms of the other. Likewise, it is a mistake to identify activation of sensory-motor regions of the brain with either perceptual or amodal symbolic processes. Research might have identified categorical reasoning as strongly *correlated* with sensory-motor regions (sect. 2.1), but this is not a sufficiently strong claim to warrant a rejection of amodalist approaches that are perfectly consistent with such evidence (other commentators, Adams & Campbell 1999; Aydede 1999; Zwann et al. 1999, have made a similar point, though with different emphasis). In his response to the commentaries, Barsalou has replied to this general argument on the grounds that amodal approaches do not fit with behavioral findings involving occlusion and size perception, and that patients showing sensory motor – but not conceptual knowledge – deficits would be frequently observed if





amodalism were true. But, again, such empirical evidence conflates correlation and causes, and it is not clear from his reply whether Barsalou realises that the burden of plausibility rests with the newcomer theory he is advancing, not the orthodoxy (in the following I shall suggest another response Barsalou can raise against amodalist objections).

I want to look at the fourth and fifth difficulties – the transduction and the symbol grounding problems. Here it seems that Barsalou really has a case. However, I shall suggest that his argument supports something far more subtle and enriched than the perceptual systems approach he advances.

Barsalou suggests that amodal symbols are arbitrarily related to the perceptual states they encode in a similar way to “how words typically have arbitrary relations to entities that produce them.” In particular, such symbols are “linked arbitrarily to the perceptual states that produce them” (p. 578). “Just as the word “chair” has no systematic similarity to physical chairs, the amodal symbol for *chair* has no systematic similarity to perceived chairs” (pp. 578–79). The word “chair” is arbitrary in nature and conventional in its genesis: we might have had another word to describe perceptual states of the chairy kind. Similarly, there is no principled reason why the amodal token that represents chairs (i.e., *chair*) needs to be the token it is, and not some other token. Hence, the problems of transduction and symbol-grounding arise for amodalist views: (1) How is the arbitrary symbol represented grounded in the transduced sensory states (how does the neurally embedded amodal expression arise from sensory impingings)? (2) How do we map the mental token *chair* to the thing in the world it represents (how does the expression map back to the chair)? The amodalist story assumes that the arbitrary symbols that do this job are structured symbolic expressions, but it is hard to see exactly how they can meet these problems without involving perceptual representation (Harnad 1990); and if they do, Barsalou’s point is that perceptual symbols are all that are needed.

Are all amodal symbols essentially arbitrary? Onomatopoeic symbols don’t seem to be. The word “creak” really does seem to represent the sound of, say a door creaking – and not in an arbitrary way. The symbol is crucially perceptual. Yet this symbol is also amodal: it is structured and proposition-like (yet grounded in the perceptual aspect of the world it represents). Suppose there were structured amodal symbols that did the same job – that is, they neurally encoded symbols that represent perceptual states in the same way as onomatopoeic symbols represent sounds. Would these face the same objections as conventional amodal symbols? It is hard to see how structured symbols such as propositions can stand in the face of the transduction and symbol grounding problems, but perhaps these objections could be overcome if it were found that a different account of symbols could be sustained.

Barsalou’s solution is to reject amodalist approaches entirely and plump for a perceptual symbol theory. These representations stand in an entirely different relation to the proximal stimulation that produced them than do amodal symbols. In particular, they stand to the thing represented as an *analogue* of the perceived entity. This process works via the medium of selective attention. Continual promptings of the associative areas of the sensory motor regions of the brain results in the perceiver being casually driven to enter certain categories they represent. Barsalou argues persuasively that this way of understanding the connection between representation and the thing represented caters to familiar features of representations such as unbounded generativity and recursive elaboration (sect. 3.1 of the target article) and so has important advantages over amodal approaches. As well, it is consistent with various kinds of connectionist approaches (sect. R5.2 of the response).

Another possibility, however, is that the distinction between modal and amodal symbol systems amounts to a false dichotomy. Suppose, instead, that the brain represented the world in a way which contained aspects of both characteristics in most cases (although there might be singular instances of *strictly* modal symbols

for, say, abstract ideas such as justice, and strictly perceptual symbols for qualia, such as pain). That is, just as it makes no sense to call a pH neutral soil acidic or alkaline, so it makes no sense to call most representations “modal” or “amodal” except at the polarities of a continuum of content. Call this a *continuum* account of representation. In this view, most day-to-day representations would be something like onomatopoeic symbols – with both perceptual and nonperceptual aspects. This way of understanding how the brain represents the world would preserve the good sense of Barsalou’s recognition of perceptual representations and yet avoid the tendency to minimise the gains won with symbolic representations so vital to contemporary cognitive science. It would also be consistent with an evolutionary account of how mental content might have been brought about (Davies 1996).

Barsalou (1999r, p. 638, sect. R1.3) admits that both modality-specific and modality-general systems may well exist. He also admits (Barsalou, personal communication) that the system he proposes contains mechanisms that go beyond perception and that rely heavily on associative areas; note his constant appeal to Damasio’s convergence zones (cf. Damasio 1989). Why not admit that a mixture of approaches may be needed in understanding representation itself? Elsewhere, Barsalou acknowledges that because selective attention is flexible, it serves the role of “establish[ing] symbols that serve higher goals of the system” (R2.2, pp. 641–42). Now, it must be wondered just what Barsalou’s “perceptual symbols” are if they are meant to bear the load of both lower end perceptual integrations and higher order goals. In what sense are they perceptual? “Perception” is being used in a very attenuated sense indeed.


An account which was both perceptual and served “higher goals” would, I think, be of interest to both Barsalou and defenders of amodalism. Only a continuum account could include such considerations. Of course, the details would need to be worked out, but the point I am making is that the deficiencies of amodalism do not necessarily support a perceptual symbols theory, but perhaps something else entirely.


A continuum account might enable Barsalou to answer his amodalist critics in the following way: Although there is certainly evidence for amodalism in the area of concept and category formation, this evidence does not necessarily mitigate against perceptual representation. Representation is more complex than hitherto imagined. Barsalou is right in pointing out that a correction is needed in the progress of amodalist views. However, he might be wrong in thinking that perceptual symbols alone will do the job. Deciding between these modes of representation assumes a false dichotomy. The real question is not: how do we decide between modal and amodal perceptual systems? The real question is: *How can representations have both perceptual and nonperceptual aspects?*

References

Letters “a” and “r” appearing before authors’ initials refer to target article and response, respectively.

- Adams, F. & Campbell, K. (1999) Modality and abstract concepts. *Behavioral and Brain Sciences* 22(4):610. [MD]
- Aydede, M. (1999) What makes perceptual symbols perceptual? *Behavioral and Brain Sciences* 22(4):610–11. [MD]
- Barsalou, L. E. (1999t) Perceptual symbol systems. *Behavioral and Brain Sciences* 22(4):577–609. [MD]
- (1999r) Perceptions of perceptual symbols. *Behavioral and Brain Sciences* 22(4):637–60. [MD]
- Damasio, A. R. (1989) Time-locked multiregional retroactivation: A systems-level proposal for the neural substrates of recall and recognition. *Cognition* 33: 25–62. [MD]
- Davies, W. M. (1996) *Experience and content: Consequences of a continuum theory*. Avebury Series in Philosophy. Ashgate, Aldershot. [MD]





Continuing Commentary

- Harnad, S. (1990) The symbol grounding problem. *Physica D* 42:335–46. [MD]
- Kosslyn, S. M. (1994) *Image and brain*. MIT Press. [MD]
- Lang, P. J. (1979) A bio-informational theory of emotional imagery. *Psychophysiology* 16:495–512. [MD]
- Shepard, R. N. & Metzler, J. (1971) Mental rotation of three-dimensional objects. *Science* 171:701–703. [MD]
- Zwann, R. A., Stanfield, R. A. & Madden, C. J. (1999) Perceptual symbols in language comprehension: Can an empirical case be made? *Behavioral and Brain Sciences* 22(4):636–37. [MD]

Lawrence W. Barcelou has declined to respond to the above continuing commentary.

Commentary on Stephen E. Palmer (1999). Color, consciousness, and the isomorphism constraint. *BBS* 22(6):923–989.

Abstract of the original article: The relations among consciousness, brain, behavior, and scientific explanation are explored in the domain of color perception. Current scientific knowledge about color similarity, color composition, dimensional structure, unique colors, and color categories is used to assess Locke's "inverted spectrum argument" about the undetectability of color transformations. A symmetry analysis of color space shows that the literal interpretation of this argument – reversing the experience of a rainbow – would not work. Three other color-to-color transformations might work, however, depending on the relevance of certain color categories. The approach is then generalized to examine behavioral detection of arbitrary differences in color experiences, leading to the formulation of a principled distinction, called the "isomorphism constraint," between what can and cannot be determined about the nature of color experience by objective behavioral means. Finally, the prospects for achieving a biologically based explanation of color experience below the level of isomorphism are considered in light of the limitations of behavioral methods. Within-subject designs using biological interventions hold the greatest promise for scientific progress on consciousness, but objective knowledge of another person's experience appears impossible. The implications of these arguments for functionalism are discussed.

Color, qualia, and psychophysical constraints on equivalence of color experience

Vincent A. Billock^a and Brian H. Tsou^b

^aGeneral Dynamics, Inc., U.S. Air Force Research Laboratory, Suite 200, 5200 Springfield Pike, Dayton, OH 45431; ^bU.S. Air Force Research Laboratory, WPAFB, OH 45431. Vince.Billock@wpafb.af.mil
Brian.Tsou@wpafb.af.mil

Abstract: It has been suggested that difficult-to-quantify differences in visual processing may prevent researchers from equating the color experience of different observers. However, spectral locations of unique hues are remarkably invariant with respect to everything other than gross differences in preretinal and photoreceptor absorptions. This suggests a stereotyping of neural color processing and leads us to posit that minor differences in observer neurophysiology may be irrelevant to color experience.

Whenever a philosopher corners a psychophysicist, the qualia problem is likely to be raised. As card-carrying members of the second camp, we have often been asked some variation of: Is my experience of (insert your favorite color) the same as yours? Our answer has generally been that equivalent color experiences are quite likely if you let us specify how the color is created. This answer is driven by the common experience of psychophysicists that color processing/experience is remarkably replicable and somewhat stereotyped (Rubin 1961; Boynton 1966),¹ and is supported by psychophysical analogues of arguments made in sections 3.3–3.4 of Palmer (1999). Palmer's excellent target article motivates a deeper analysis of the constraints that color psychophysics imposes on equating color experiences.

Palmer points out that if neural activity is identical, it is unparliamentary to posit a difference in color experience. Conversely, Palmer argues that the multitude of large and small cortical differences between observers makes the decision about an exact neural match problematical. There is, however, another approach that employs psychophysical performance linked to a neural correlate. In color opponent theory, unique green, blue, and yellow are considered the null points of opponent (usually subtractive) operations between mechanisms driven by L-, M-, and S-cone photoreceptors. As such, the unique hues provide a strong constraint on the specification of the two independent red-green and blue-yellow color opponent channels. Similarly, the spectral locations of balanced orange or cyan constrain the relative scaling of

the two channels. If two subjects share the same unique hue, then we know that they have identical (zero) neural responses in the nulled channel. Moreover, if their balanced hues are also the same, then we know that the un-nulled channel is scaled the same in both observers and that the neural responses in these channels are also nearly identical. So, for example, if two observers have the same unique yellow and the same balanced orange, then when the monochromator is set to the unique yellow point, both observers experience the same responses in their color opponent channels: zero in the red-green channel and a yellow response in the blue-yellow channel that is tightly constrained by the identicalness of the balanced orange setting. In this context, note that we do not train the observers (or rely on society training them) to see particular colors only within a few nanometers range; we just ask them to use their color system as a nulling instrument – like a Wheatstone bridge – something subjects are extremely good at (Hurvich & Jameson 1974; Regan 1991). Of course, given the concepts of metamerism and stimulus equivalency, it is unnecessary to restrict our analysis to identical neural responses to identical stimuli, but doing so facilitates making a second point on the stereotyped nature of the neural processing of wavelength.

Consider Rubin's (1961) study of unique hues in color normals and anomalous trichromates with decent color discrimination. Rubin asked 278 color normals (determined by Rayleigh match) to use a monochromator to scan the spectrum and find the three unique spectral hues (unique green, blue, and yellow) and two balanced hues (balanced orange and cyan). Subjects were instructed, for example, to find the yellow wavelength that contained no trace of red or green, or to find the orange wavelength that contained equal amounts of red and yellow. Rubin used a bracketing procedure to eliminate the spectral order effects that would otherwise occur (Beegan et al. 1999). Rubin found that normal subjects all selected wavelengths within a few nanometers (nm) of each other. Subjects with abnormal L-cone (protanomals) or M-cone pigments (deuteranomals) also clumped together, with results similar to normals, but shifted in the direction expected by color theory (see Table 1).^{2,3}

The distribution of wavelength settings for a given color is very tight (SD of 2 nm for yellow and barely worse than the within-subject test-retest variability). And, today the narrow distributions that Rubin measured could likely be tightened by genetic screening for minor variations in photopigment maxima and optical screening for excessive preretinal absorptions (which may contribute to observer variability, but seem to have minor effects).³

