action. This formulation is congenial to the P&V project, but is less extreme, because it allows for unconscious representation: R represents object O if, because of some appropriate isomorphism, R plays the role of O in a simulated action involving O, either consciously or preconsciously (e.g., habituatedly). Given this possible formulation, the P&V exclusion of nonconscious representation seems unnecessarily counterintuitive.

In standard computational views, unconscious representations undergo transformations, resulting in behavioral outputs. These transformations of the unconscious representations are what constitute the unconscious computations. Consciousness is merely an optional way to access the results of the computations. For many computationalists, this renders consciousness an epiphenomenon, whose only causal powers over behavior or thought would have to be illusory (Jackendoff 1987).

Even those convinced that consciousness is not merely the epiphenomenon of information processing, but also requires processing in the unique manner of an active, self-organizational system, should notice that P&V force a choice between extreme viewpoints and ignore much middle ground. One successful research program frequently touted as perfectly compatible with computationalism involves different layers of sensory cortex in occipital and temporal lobes performing computations on incoming perceptual signals (Hubel & Wiesel 1959; Richardson 1991). These transformations are obviously unconscious, because when there is virtually complete occipital activation in response to a completely unexpected stimulus (indicating that the transformations are virtually complete), the subject still lacks perceptual consciousness unless there also occurs a parietal 300P electrical potential (Aurell 1983; 1984; 1989; McHugh & Bahill 1985; Srebro 1985; Weiskrantz 1986). Occipital and temporal lobes can do everything they normally do in processing the perceptual data, including the 100P occipital potential and the 200N "mismatch negativity," without the subject having consciousness of the stimulus. (In ERPs, the numbers refer to milliseconds after presentation of the stimulus. Extensive processing occurs during the first 250 msec of processing, with or without the consciousness accompanying the 300P.) These unconscious occipital transformations fit a computational paradigm: Cells in consecutive layers of sensory cortices analyze different features of perceived objects – lines, angles, shapes, colors, and so forth. These sequences of transformations are unimaginable on a conscious basis; we cannot imagine consciousness of color without shape or vice versa, yet our sensory cortices "compute" these properties separately and then recombine them.

Notice that the occipital transformations of perceptual signals are used to explain how a certain type of representation *comes about* in the first place. In one sense, we think of the pre-occipital signal (as received in the thalamus, for example) as an "unconscious representation," which will then be combined with other signals and transformed into a more fully developed "representation" – a representation in a different sense. But this highlights the need to think more carefully about what constitutes a "representation": Newton's enactive, self-organizational view of representation would eliminate the problems just mentioned by treating occipital activities as "potential" representations – activities that will lend themselves to use by the organism in representational action if the occasion should arise. The same could be said for representations in memory, thought, unconscious emotions, and the like. Representations and computations can occur without all conscious processes reducing to them.

Natura non facit saltum: The need for the full continuum of mental representations

Robert M. French
Department of Cognitive Science, University of Liège, Liège, Belgium.
rfrench@ulg.ac.be

Abstract: *Natura non facit saltum* (Nature does not make leaps) was the lovely aphorism on which Darwin based his work on evolution. It applies as much to the formation of mental representations as to the formation of species, and therein lies our major disagreement with the SOC model proposed by Perruchet & Vinter.

Perruchet & Vinter (P&V) admit, of course, that conscious representations emerge from an underlying neural substrate. But the type of emergence for which they argue seems to involve a sudden, quantal leap from the unconscious to the conscious. One moment, the representation of an object, a scene, or a situation is in the process of being generated and is of no importance whatsoever in any cognitive calculus; and then, suddenly, as if by magic, the representation bursts into consciousness, thereby becoming endowed with all the cognitive powers of conscious representations. P&V write:

> mental life comprises only two categories of events: The conscious representations and the unconscious processes generating those representations. The two are linked like the head and tail of a coin... [the] processes and mechanisms responsible for the elaboration of knowledge are intrinsically unconscious, and the resulting mental representations and knowledge are intrinsically conscious. (target article, sect. 1.3.1)

It strikes us that a gradualist picture of representation-formation – for example, the classical Hebbian cell-assembly framework will suffice – would, in one fell swoop, explain most, if not all, of the instances of unconscious influences on conscious processing that the authors work so hard to explain away within their no-unconscious-representations SOC framework. Further, this new framework would in no way undermine the associationist principles that drive their model (correctly, in our opinion). Accepting the existence of representations that run the gamut from the embryonic and unconscious to the fully formed and conscious in no way implies the need for a "sophisticated unconscious processor."

The authors accept the notion of graded and partial conscious representations. Within a simple neural network framework, there is no problem extending these notions to unconscious representations. As it stands, the authors would have a great deal of difficulty in their SOC framework in distinguishing between an "unconscious representation" and an "absent representation" (i.e., no representation), because they would maintain that both situations have no effect whatsoever on conscious perception. But, as we hope to show in the thought-experiment presented here, there must be a difference. This difference, if a real-world version of the thought-experiment were actually run, would presumably be able to be measured with appropriately sensitive instruments.

Suppose that two individuals, A and B, start with perfectly identical brains. Via a rigid, completely reproducible procedure, A learns the concept Ω, and B does not. Now, presumably, learning Ω involves a physical (presumably, synaptic) modification of a specific set, S_A, of neurons in A's brain. The precisely corresponding set of neurons in B's brain, S_B, undergoes no such physical change. Presumably, P&V would say that the concept Ω is now physically represented in A's brain (whether active or not). Now, since they explicitly accept the concept of representational decay, we will suppose that the synaptic changes that constituted A's representation of Ω gradually decay in precisely the reverse order in which they were strengthened when A was originally learning the concept Ω. Further suppose that we have a device capable of stimulating the neurons in S_A (and only those neurons). At some point during this decay toward the original state of the neurons before A learned Ω, A would presumably no longer be consciously aware of the concept Ω when S_A was stimulated. (This point will be somewhere in the zone corresponding to A's very early learning of the concept, before the representation would be conscious. P&V explicitly concede that there is such a period.) At this point, we now have A relearn Ω and B learn Ω for the first time, employing exactly the same procedure originally used when A first learned Ω. Surely, P&V would agree that A would relearn the concept Ω faster than B because, as we have set things up, A will have a representational "head-start" over B. We thus have a very simple hypothetical case of how an unconscious representation could significantly affect the conscious experience of concept acquisition. Further, the SOC account, relying as it does only on conscious representations, would be at a loss in explaining this learning-time difference, unless they took the unfalsifiable position that A's more rapid learning of Ω simply demonstrated that the decayed representation with which A started prior to relearning Ω must, in fact, have been conscious all along.

It may well be that there is, indeed, some sort of "connectivity phase change" when a neural representation has the possibility of becoming conscious when activated. This could be the point described by Hebb as when "reverberation in the structure might be possible . . . reverberation which might frequently last for periods of time as great as half a second or a second, [this being] the best estimate one can make of the duration of a single 'conscious content'" (Hebb 1949, p. 74). But if one is to present a coherent picture of cognition that takes into account neural, representational, and cognitive phenomena, one must not neglect the representational stages leading up to this creation of cell-assemblies or, in the language of P&V, up to the emergence of fully conscious representations.

In conclusion, we suggest that the SOC model might do well to turn to basic neural network principles that would allow it, without difficulty, to encompass unconscious representations, as described above. (See, e.g., Cleeremans & Jiménez 2002; Mathis & Mozer 1996.) These "unconscious" representations – some of which may evolve into representations that, when activated, would be conscious – can affect consciousness processing, but do so via the same basic associative, excitatory, and inhibitory mechanisms that we observe in conscious representations. The inclusion of this type of representation in no way requires the authors to also posit sophisticated unconscious computational mechanisms.

Unconscious semantic access: A case against a hyperpowerful unconscious

Daniel Holender[a] and Katia Duscherer[b]

[a]*Laboratoire de Psychologie Expérimentale, Université Libre de Bruxelles 50, Brussels, Belgium;* [b]*Laboratoire de Psychologie Expérimentale, Université René Descartes (Paris V), Boulogne-Billancourt, France.*
holender@ulb.ac.be kduscher@ulb.ac.be

Abstract: We analyze some of the recent evidence for unconscious semantic access stemming from tasks that, although based on a priming procedure, generate semantic congruity effects because of response competition, not semantic priming effects. We argue that such effects cannot occur without at least some glimpses of awareness about the identity and the meaning of a significant proportion of the primes.

Like Perruchet & Vinter (P&V), we fully endorse a mentalistic perspective, which implies that we do not posit the existence of a "powerful," or more precisely, an *intentional* cognitive unconscious. Thus, we basically share the view of Searle (1990; 1992) and Dulany (1997) that the unconscious is intentional in a dispositional way. In this commentary, we expand on the claim made by P&V in section 8.2 that the available data on unconscious semantic access do not constitute a challenge to the mentalistic framework.

In assessing the plausibility of the evidence for unconscious semantic access, a distinction must be made between tasks generating semantic priming effects and tasks generating other effects based on stimulus meaning, such as Stroop and Stroop-like congruity effects. This distinction has been somewhat blurred in recent work, maybe partly because of the multiple meanings of the term *priming*, which can designate an experimental procedure, an observed effect, and a hypothetical causal process, such as automatic spreading activation in semantic memory (e.g., Neely 1991). Much of the early evidence for unconscious semantic access under masking, criticized by Holender (1986), was based on a semantic priming paradigm yielding bona fide semantic priming effects. Much of the recent evidence for unconscious semantic access discussed by P&V does not qualify as priming because it rests on tasks that, although based on a priming procedure, are functionally equivalent to Stroop-like tasks. These tasks are generally assumed to generate congruity effects because of response competition (e.g., Eriksen 1995; Holender 1992; MacLeod 1991), not priming effects.

The studies of Greenwald et al. (1996; Draine & Greenwald 1998) are based on prime and target words with strong positive and negative affective connotations. The SOA between the prime and the target is very brief (under 100 msec), and the prime is interleaved between two masks consisting of random letters strings. Even though the primes could not be discriminated above chance, the binary classification of the target words in terms of their pleasantness is more accurate in congruent trials, in which the polarity of the prime and the target words are the same, than in incongruent trials, in which the polarities are opposite. Similarly, in the studies of Dehaene et al. (1998; Naccache & Dehaene 2001), which are based on a comparable procedure, the speed of classification of a single-digit target number in terms of whether it is larger or smaller than five is affected by the congruency of the unconscious prime number.

Initially, Greenwald et al. (1996; Draine & Greenwald 1998) interpreted their finding as reflecting semantic priming based on spreading activation. Then, Klinger et al. (2000) demonstrated that this effect does not depend at all on spreading activation but on response competition. This was taken as evidence that the unconscious primes must be covertly classified according to the same rule as the one applied to the visible target (see also Dehaene et al. 1998). Next, it was shown that the congruity effect only appears with primes that have been used repeatedly as targets (Abrams & Greenwald 2000; Damian 2001), which prompted a reinterpretation of the effects in terms of the formation through learning of a direct stimulus-response link based on superficial features of the stimuli. However, Abrams et al. (2002) argued that this link must rather be established between the stimuli and the semantic categories, as the learning effect resisted a change in response assignment. Nevertheless, Naccache and Dehaene (2001) persisted in their account in terms of unconscious semantic classification, because the congruity effect still occurs with unconscious primes, which have not been seen before as targets.

All these interpretations of unconscious congruity effects rest on the assumption that the primes are completely unavailable to awareness. If correct, they imply a hyperpowerful unconscious, that is, an unconscious even more powerful than the one already required to explain unconscious semantic priming effects. We contend that this conception is profoundly mistaken because, as was pointed out by Prinz (1997), a stimulus has no inherent information sufficient to specify a response outside the context of a goal-directed task imposed by the instructions. Besides, the primary source of response conflicts underlying the congruity effects described above must lie in conscious mental representations (cf. Holender 1992), because there is no stored information, and hence no information that can be automatically activated, about whether a number is smaller or larger than five or about whether the concept denoted by a word has a pleasant or unpleasant connotation. Therefore, the only possible source of conflict lies in the fact that most participants think about the irrelevant information in terms similar to those used by the instructions to describe how