very constrained by our linguistic knowledge (we replay the signal as we decoded it; we cannot look at it again), unless the words or the strings of words are repeated numerous times in the same context, which is unlikely in real-life situations. One way to circumvent this limitation is through the sensori-motor loop of language production, but children do not repeat or produce all the words that they understand, so this is obviously not the only route into language structure.

However, we think that these limitations do not invalidate PARSER and the principle of SOC, if one takes into account the differences between the oral modality and other modalities. For other modalities, data are often not transient and it is possible to take full advantage of SOC to use the world as an "outside memory" and to trade representation against computation. Visual presentations and tactile sensations are much more likely than oral stimuli to lead to complex representations. Now, what is true for oral language is not true for written language. Thanks to the visual format of written language, it becomes much easier to build conscious representations isomorphic to the outside world. What necessitated thousands of syllables for oral language (Perruchet & Vinter 1998b) can be accomplished with much fewer repetitions. This makes it easier to analyze language per se, further disconnected from its semantic interpretation, than it was for oral language.

Our proposal is that SOC works with both oral and written language but results in different products, and that structural differences between these products have great impact on the development of language in the child and on the nature and structure of language.

The characteristics of oral language make it difficult for young children to analyze the oral signal into parts, although they are already able to understand and manipulate a great number of words. Children actually take quite a long time before being able to split into parts and manipulate the oral forms that they used as unanalyzed wholes (Peters 1983) and exhibit knowledge of syntactic categories (Tomasello 2000a). However, the greater permanence of visual and tactile representations allows children to use these representations as a support for their first oral stimuli representations. This is consonant with their item-based linguistic behavior (Tomasello 2000b) and with the importance of the semantic nature of their first linguistic knowledge. Following the preschool years, the structure of written language becomes the outside support for internal isomorphic representations, which become in turn a support for reasoning about language. Indeed, metalinguistic thinking appears when children have learned to read, not before (Gombert 1992).

As representations can be much more complex with written material than with oral material, real or seemingly rule-governed behavior would be much more developed in written language than in oral language. This can explain a lot about the differences between oral and written language. These differences can be found in the grammatical structures produced by native speakers (Miller & Weinert 1998), or in the structures and interpretations of texts and stories (Ong 1982). Oral language tends to be more formulaic than written language and mostly composed of first-order structures. Also, oral language structures and representations are more semantics-centered than syntax-centered with a lot of sometimes very complex, item-based constructions (see Croft 2001), as is the case for preschool children (Tomasello 2000b).

Finally, in the history of the human civilizations, the metaanalysis of language always appeared after the discovery of writing (Auroux 1994). The first written productions always were ideographic and only later became phonologic through a lengthy historical process, which eventually led to the development of the linguistic sciences (see Auroux 1994; Harris 1980). It seems as if a visual support was a necessary prerequisite in the detailed elaboration of the structure of the oral support and help in the transition from a purely semantic-centered view of language to a purely syntactic-centered view.

Our conclusion follows the proposal of Auroux (1994), that there are different levels of language consciousness: first, epilinguistic knowledge (the unconscious knowledge of language), which is a pure product of oral language; and, second, metalinguistic knowledge (the conscious knowledge of language), which is a byproduct of the use of written language. This opposition is unfortunately not yet taken into account in most linguistic and psycholinguistic theories, but it could adequately be described and explained within the SOC framework.

Neo-associativism: Limited learning transfer without binding symbol representations

Steven Phillips

Neuroscience Research Institute, National Institute of Advanced Industrial Science and Technology, Tsukuba, Ibaraki, Japan. **steve@ni.aist.go.jp**
http://www.etl.go.jp/~stevep

Abstract: Perruchet & Vinter claim that with the additional capacity to determine whether two arbitrary stimuli are the same or different, their association-based PARSER model is sufficient to account for learning transfer. This claim overstates the generalization capacity of perceptual versus nonperceptual (symbolic) relational processes. An example shows why some types of learning transfer also require the capacity to bind arbitrary representations to nonperceptual relational symbols.

Perruchet & Vinter (P&V) claim that with the addition of a simple principle, the capacity to determine whether any two arbitrary stimuli are the same or different, their association-based PARSER architecture is sufficient to explain learning transfer, such as that reported by Marcus et al. (1999). Learning transfer relies on between-stimulus relations that are common across different sets of unique stimuli. However, the concept of a common relation belies an important distinction between perceptual and nonperceptual (symbolic) relations that has different implications for generalization capacity. Consequently, although their proposal may be sufficient to explain, for example, transfer in infants based on common perceptual relations; it is not sufficient for transfer based on common nonperceptual relations. And, by extension, that infants demonstrate transfer on a phoneme sequence recognition task, does not imply the capacity to process symbolic relations.

Perceptual versus nonperceptual (symbolic) relations. A perceptual relation is computed from the perceived features of the related stimuli. *Sameness/difference* is a perceptual relation where there exists a thresholded distance function over any two points in feature space. Computing perceptual relations affords learning transfer in relation-match-to-sample, because the task is reduced to learning that a match is found when the distances between paired stimuli in source and target pairs are either both less than, or both greater than threshold. Because matching perceptual sameness/difference relations is the essence of the tasks analyzed, I agree with the authors that PARSER is capable of transfer on these tasks. By contrast, a nonperceptual (or, symbolic) relation, such as *sibling*, is one where no such function over perceptual features is available.

Of itself, this distinction is unremarkable. But, in the context of characterizing cognitive processes it lies at the root of the problem which appears in the authors' claims and the sorts of data they are attempting to explain. Perceptual relations can be *imposed* on nonperceptual relations by prior learning of new perceptual relations, and these learned perceptual relations can be used as the basis for transfer (e.g., Hinton 1990). The problem for the authors (and this type of explanation, generally) is how much of this prior learning justifiably explains learning transfer. It seems reasonable to suppose that infants already have the capacity to distinguish arbitrary pairs of phonemes prior to the phoneme sequence recognition task. But, for other tasks, such as Hinton's Family Tree and Transverse Patterning, no amount of prior learning justifies learning transfer in the networks analyzed (Phillips 1999; 2000). If computing sameness/difference perceptual relations is sufficient, how

does one explain transfer when this sort of information is not available? A nonperceptual (symbol) version of relation-match-to-sample is defined to illustrate the problem and the limits of models sensitive only to perceptual relations, including PARSER.

Symbol-relation-match-to-sample. In this task, the subject is presented with four pairs of stimuli: AB, CD, EF, and GH. Symbol-relation-match-to-sample differs from relation-match-to-sample in that the sameness/difference relation is specified by colored backgrounds, not features specific to the stimuli. Here, brackets indicate colors, so that {A}{B} – identical background colors – identifies the symbolic relation SAME (A, B); and {A} [B] – different background colors – identifies the symbolic relation DIFFERENT (A, B). The subject is then presented with an uncolored probe pair and the three remaining uncolored choice pairs. Subjects must select the choice pair with the same symbolic relation as the probe pair. For example, given pairs {A}{B}, {C}{D}, {E}[F], and {G}[H]; uncolored probe pair AB; and uncolored choice pairs CD, EF, and GH: The correct response is CD. For probe EF, the correct response is GH.

The critical aspect of this task is that the correct response cannot be based on any perceptual relation between the uncolored paired stimuli, because all uncolored stimuli are different, and any coincidental perceptual relation used to predict the response in one trial would fail on a subsequent trial where the pair were reassigned to a different symbolic relation. Thus, no amount of prior learning on perceptual relations can account for transfer on this task. Transfer is achieved with the capacity to bind arbitrary pairs of stimuli to symbols (representing the symbolic relations SAME and DIFFERENT); and depending on the decision process, either the capacity to match retrieved pairs from relational symbols or match retrieved relational symbols from pairs.

Marcus (1998) has argued that cognitive processes must be symbolic, because only symbol systems can explain the capacity to generalize across novel stimuli. Yet, the infant data used as evidence of symbol systems are open to two different interpretations. This comment is not to deny the importance of the authors' model or Marcus et al.'s infant data. To the contrary, they have helped us formulate a tripartite theory of cognitive processes, where this type of generalization behavior lies at an intermediate level between symbolically structured and nonstructured elemental associative processes (Halford et al., submitted).

Associative learning: A generalisation too far

Martin Redington

Department of Psychology, University College London, London, United Kingdom. **m.redington@ucl.ac.uk**
http://www.homepages.ucl.ac.uk/~ucjtfmr/

Abstract: I argue that Perruchet & Vinter's claim that representations are conscious, and processes unconscious, gives too much ground to the cognitive unconscious; and that the boundary between conscious and unconscious mental phenomena is unlikely to fall neatly along these lines. I also propose that in the absence of more detailed models that demonstrably provide a reasonable account of the data, claims that associative mechanisms may underlie all cognition are premature.

I am greatly sympathetic to the general thrust of the argument advanced by Perruchet & Vinter (P&V) and in close agreement with their interpretation of much of the evidence that they cite.

For implicit learning, associative mechanisms provide the best available account of human performance, including the transfer phenomenon. There is no convincing evidence for rule learning, or for the existence of unconscious knowledge.

For language acquisition, structural information in the input, accessible to associative learning mechanisms, contains a great deal of information about linguistic structure, which is potentially available to the learner.

Given the success of associative learning mechanisms in these two domains, one may reasonably speculate, as P&V have done, that similar mechanisms may underlie many other aspects of cognition.

However, P&V's thesis goes far beyond this, making admirably bold claims about the accessibility of representations and processes to conscious awareness, and on the nature of learning mechanisms underlying human cognition generally. I believe that these generalisations are flawed.

First, consider the proposal that all representations are consciously accessible and the processes that generate them are unconscious (where representations are mental events that "assume the function of some meaningful component of the world" [target article, sect. 1.3.2] such as a person or object).

Although computational models have generally neglected issues of conscious accessibility, many models can be read as compatible with P&V's claim, with attention ensuring that the appropriate input is presented to a mechanism whose inner workings are inaccessible to consciousness, but whose outputs are available for further processing or action. This is especially true of connectionist models, where representations (patterns of activations over units) are transient, as in P&V's account, occurring only in response to appropriate inputs.

This compatibility suggests that P&V's claim may be a useful characterisation of some aspects of cognition, but as a general account, ironically, their position surrenders far too much ground to the cognitive unconscious. If some processing is conscious, as P&V admit, then, given their argument on representations (some representations are clearly conscious, so it is more parsimonious to assume that are no separate unconscious representations), why assume a separate class of nondeliberative and wholly unconscious processing? This assumption conflicts with Shanks and St. John's (1994) view, that in many domains, and especially implicit learning, knowledge acquisition is the result of conscious processes acting on conscious representations. While I share the skepticism of both sets of authors about many claims for the cognitive unconscious, there seem to be many mental events to which we may in principle have access but generally do not attend to, and others (e.g., neural events) to which we uncontroversially have no conscious access. I do not see a strong case for supposing that the boundary separating accessible and inaccessible mental phenomena falls neatly between representation and processing.

Second, P&V claim that associative learning and memory mechanisms are sufficient to account for human cognition, and that rule-based mechanisms are not required. I consider this separately, because one can imagine variants of P&V's account that argue that representations are conscious but permit them to be symbolic, or conversely, that rely only on associative processes, while admitting unconscious representations or knowledge.

P&V's claim is compatible with the evidence from implicit learning, but is much more problematic when applied to language acquisition. Undoubtedly, structural properties of language contain a lot of information about some aspects of linguistic structure, despite influential, but mistaken, a priori dismissals (e.g., Pinker 1984). For example, in Redington et al.'s (1998) analysis of child-directed language, word co-occurrence statistics were highly correlated with a word's syntactic category. The availability of information in the input is no guarantee that it is utilised by the learner, but in this case the cues are so informative, and so readily accessible to simple learning mechanisms, that it would be very surprising if learners did not exploit them in some way.

Nevertheless, distributional properties are far from perfectly informative about syntactic category, and it seems likely that additional cues, such as semantic and phonological constraints, prosodic information, and innate knowledge (for example, of the number and/or relative proportions of words of each category), must also contribute. While one may speculate, as P&V do, that associative mechanisms may also be able to exploit semantics, prosody, or phonology (leaving aside innate constraints), and to appropriately combine this information to identify syntactic cate-