Leiser, D. (1997). The dynamics of cumulative knowledge. Behavioral and Brain Sciences, 20(1), 76-77

The dynamics of cumulative knowledge

David Leiser
Department of Behavioral Sciences, Ben-Gurion University of the Negev, Beer Sheva, Israel. **dleiser@bgumail.bgu.ac.il**
www.bgu.ac.il/beh/leiser.html

Abstract: Exploiting existing representations implies tapping an enormous domain, coextensive with human understanding and knowledge, and endowed with its own dynamics of piecewise and cumulative learning. The worth of Clark & Thornton's proposal depends on the relative importance of this dynamics and of the bottom-up mechanism they come to complement. Radical restructuring of theories and patterns of retrieval from long-term memory are discussed in the context of such an evaluation.

Clark & Thornton's (C&T's) basic point is akin to the observation that finding optimal coefficients for multivariate linear regression is straightforward, whereas devising an appropriate nonlinear model is much harder. To do so effectively, existing knowledge must be exploited. While it is difficult to take issue with this, the full significance of the observation must be realized: no mere tactical recommendation, it implies tapping an enormous domain, coextensive with human understanding and knowledge; a domain, moreover, endowed with its own epigenetic dynamics of piecewise and cumulative learning. If it turns out that most of the novelty required to solve a type-2 problem is ultimately based on that dynamics, then it is misleading to present the enjoyment of the fruits of quite different learning mechanisms as a trick to complement uninformed learning.

Exploiting existing knowledge. C&T contend that a purely bottom-up approach is not enough, and that solving type-2 problems requires the exploitation of previous learning. This could mean transferring knowledge across connectionist systems, a notion of great potential importance, but still in its infancy. Instead, they present a range of dissimilar approaches, all considered equivalent from their very abstract viewpoint: to maximize the role of achieved representation.

But how does that knowledge develop? It is hardly itself a product of type-1 bottom-up learning. Constructions at the conceptual and theoretical level obey coherence and coordination requirements (Thagard 1989) and assume a much richer representational language. The idea was first propounded by Piaget (1974) under the name of reflective abstraction and cognitive phenocopy: what is acquired at one level must be reconstructed and reorganized at another. Piaget's views are being rediscovered in various contexts. According to Keil (1994), there is no obvious mapping from the kinds of environmental regularities most salient to simple associative systems operating on perceptual primitives to sets of explanatory beliefs. This discrepancy arises for a range of reasons. One is the richer representational language. Another may stem from restructuring. Consider cascade correlation: successively generated hidden variables are akin to epicycles; they only explain residual variance (Schultz et al.). Theories, however, may change radically. Some forms of learning may give rise to a weak restructuring, involving the accumulation of new facts and the formation of new relations between existing concepts. Others involve a radical restructuring that includes changes in core concepts, structure, and the range of phenomena to be explained (Carey & Spelke 1994; Gold 1987; Vosniadou & Brewer 1987). This phenomenon is not necessarily incompatible with Clark and Thornton's view, since no actual restructuring needs to take place. One possible interpretation is that the new conceptual structure grows alongside the old one, and eventually comes to replace it (Chi 1992); this interpretation could be handled by a modular connectionism. My point here was to indicate the types of phenomena that may occur at the conceptual level because of the latter's mode of development would be very different from developments that might take place without it.

Retrieval propensities. Another point to consider is organization and retrievability in long-term memory. Here, too, there is ample room for cumulative progress that would then dominate bottom-up factors. As experience accrues, the store of potentially relevant information grows and must do so in a way that – with an appropriate retrieval mechanism – might guide novel problem solving. Research on analogy suggests that retrieval relies mostly on superficial features and hence tends to stay within domain boundaries. Only as subjects become experts do they develop the abstract tools that enable them to represent the essence of new problems and to apply relevant schemes learned in the past. This suggests that bottom-up and local learning is superseded by more conceptually based general understanding, but only for domain experts (Gentner et al. 1993; Goswami 1991; Reeves 1994). It seems accordingly that the pattern of retrieval propensities from long-term memory does not contribute much to bottom-up learning.

In sum, I have tried to illustrate how the proposal by Clark & Thornton should be critically evaluated. Their basic point is almost trivially true. Assessing its significance demands an appraisal of the relative importance of other mechanisms involved. In the event, I have concluded that the complexities of theory formation should be acknowledged, that restructuring is not necessarily a problem, and that the structure of retrieval from long-term memory does not raise any particular difficulty except in experts.

Extracting higher-level relationships in connectionist models

Gary F. Marcus
Department of Psychology, University of Massachusetts, Amherst, MA 01003. **marcus@psych.umass.edu;**
www-unix.oit.umass.edu/-marcus

Abstract: Connectionist networks excel at extracting statistical regularities but have trouble extracting higher-order relationships. Clark & Thornton suggest that a solution to this problem might come from Elman (1993), but I argue that the success of Elman's single recurrent network is illusory, and show that it cannot in fact represent abstract relationships that can be generalized to novel instances, undermining Clark & Thornton's key arguments.

Clark & Thornton (C&T) provide a compelling argument that statistical learners face problems in learning higher-level relationships. Less compelling, though, is their proposed solution, "achieved representational spaces." The centerpiece of their argument is the incremental version of Elman's (1993) model, which they and Elman argue is learning something about grammatical dependencies like subject-verb agreement.

But has Elman's model truly abstracted a higher-order regularity that is not attested in the input? Elman himself does not show this, since the only quantitative comparison he provides is "the degree to which the network's predictions match the . . . probability distributions of the training data."

In fact, Elman's model depends entirely on the statistics of lexical concurrence, never deriving abstract higher-order relations. I discovered this by conducting network simulations that contrasted statistical and relational information (Marcus 1996a). For example, in one experiment, I trained a version of Elman's model on sentences constructed from the grammar *an X is an X*, using a set of twenty different instances of *X* (rose, duck, iguana, butterfly, etc.). What happens if we test how this generalization is applied to a novel word? Given the sentence fragment *a dax is a* ——, humans readily predict the continuation *dax*.

Elman's model behaves quite differently: while it easily learns all of the training sentences, it is unable to extend the abstract underlying relationship to the novel word *dax*. This failure is robust, unaffected by the number of hidden units, the number of hidden layers, the number of training examples, or the sequence in which those training examples is presented.

The reason the network fails to extend the abstract relations lies in its roots as a statistical approximator: within the training corpus