Does the evolutionary perspective offer more than constraints?

Wolfgang Schleidt
Institut fur Vergleichende Verhaltensforschung, Savoyenstrasse 1A, A-1160 Vienna, Austria
Electronic mail: *meli@lezvax.oeaw.ac.at*

I am embarrassed to criticize such a beautiful piece of science. Newell's approach is extremely interesting, especially his concern for the time domain of cognitive processes. His contribution to cognitive science is undisputed. SOAR, as a computer language, is a working model of a cognitive system, and it apparently works very well in areas of human cognition that involve verbal processes. In fact, since verbal processes are so important in our daily lives as well as in cognitive science, it is tempting to see them as the central processes of human cognition. But not all of human cognition is verbal (think of such sensuous issues of cognitive science as music, sex, or at the interface of our senses, the processing of olfactory or spatial information). And human verbal processing is a relatively new ability in the evolutionary time frame (one can view it as a spectacular icing on the cake of mammalian cognition). In a nut shell, I feel that a unified theory of cognition must deal with the very core of human cognition and not be restricted to the verbal domain.

Why are evolutionary issues so frequently overlooked, misunderstood, or underrated? In Newell's book, for example, evolutionary considerations take not much more than 1% of the available space. I think that this bias, at the expense of understanding the very core of our existence, stems from a reductionistic view, assigning DNA a higher degree of evolutionary relevance than the structure of the human brain (safely hidden within the skull). It basically ignores the possible evolutionary significance of the specifically human "overt behavior" (including human speech) that led to the evolution of the human brain and the human genome as we know it. This is as if we were to assign to the human DNA a higher degree of scientific validity than to the specifically human features of the human brain, and as if these hidden features of the human brain were "more real" than the specifically human features of the overt human body and overt human behavior.

In Newell's *Unified theory of cognition* the bias toward biology on the lowest level only finds its expression, for example, in the label "The Biological Band" (p. 123) for the lowest part of the time scale of human action (below msec) – as if all other bands of human action were nonbiological. Why should the top of the time scale of human action be less biological? Take, for example, the cognitive problems involved in the processing of human faces and facial expressions. An essential part of the cognitive architecture for facial recognition in mammals (including man) has been localized in the temporal lobe (e.g., Rolls et al. 1985) and it is reasonable to assume that human facial recognition involves a similar cognitive architecture in this general area of the brain. Human facial expression is processed within the time scale of the "cognitive band" and subject to a striking illusion (Schleidt 1988) within this time frame; I think it is most reasonable to assume that this "biological" illusion involves processes ranging from the "biological band" well into the range of the "social band" (days to months).

Processing facial information is only one portion of the "biological roots" of human sociality, however; common obstacles to understanding the biological underpinnings of our existence are ignorance and biases we help propagate, since we cannot struggle free of our own culture. This may sound terribly "biologistic," and what I mean may be better understood by nonbiologists when worded by an eloquent linguist.

Benjamin Lee Whorf, arguing against "armchair generalizations about grammar, and the related fields of logic and thought-psychology" based solely on "Standard Average European" languages, made the following observation, which speaks equally well against generalizations about cognition based solely on verbal features:

> The evolutionary concept, having been dumped upon modern man while his notions of language and thought were based on knowledge of only a few types out of the hundreds of very diverse linguistic types existing, has abetted his provincial linguistic prejudices and fostered the grandiose hokum that this type of thinking and the few European tongues on which it is based represent the culmination and flower of the evolution of language! This is as if a pre-Linnaean botanist who had conceived the idea of evolution should suppose that our cultivated wheat and oats represent a higher evolutionary stage than a rare aster restricted to a few sites in the Himalayas. From the standpoint of a matured biology, it is precisely the rare aster which has the better claim to high evolutionary eminence; the wheat owes its ubiquity and prestige merely to human economics and history. (Whorf 1956, p. 84)

This paragraph was apparently written in the late thirties, and "from the standpoint of a matured biology" half a century later, both the rare aster and our cultivated wheat are seen as evolved equally "highly"; all species surviving up to this moment are adapted to their environment. Nonverbal skills may well be "the rare aster," and verbal skills "the wheat." But even if we accept human verbal skills as "the culmination and flower of the evolution" of cognition they appear too specialized to serve as the core of cognition in general.

The search for *one* unified theory was initiated by physicists and an image comes to my mind in this context. Picture for yourself what a pleasure it would have been for Archimedes to be allowed to look into the future and see the beauty of Swiss clockwork. He could rightfully take pride in his ingenious contribution to classical physics. But would you trust Archimedes to build a unified theory of physics, based solely on his discoveries of the principle of the lever and without reference to Einsteinian relativity?

Choosing a unifying theory for cognitive development

Thomas R. Shultz
Department of Psychology, McGill University, Montreal, Quebec, Canada H3A 1B1
Electronic mail: *shultz@hebb.psych.mcgill.ca*

I am entirely sympathetic to Newell's argument that the time is now right for beginning to generate unified theories of cognition. My commentary focuses on the question of whether his proposed SOAR model is the right vehicle for such a unified theory, particularly for cognitive developmental phenomena.

A first glance suggests that SOAR has much to recommend it as a candidate unified cognitive theory. It provides a remarkably smooth integration of learning, knowledge, and problem solving. Although condition-action rules have long been recognized, largely due to Newell's pioneering efforts, as a convenient and natural way to represent human knowledge, the issue of how rules are learned has been difficult to resolve. The solution in SOAR, chunking the results of search-based problem solving, is far more elegant and computationally tractable than trying to coordinate the various other learning mechanisms that have been proliferating in the literature.

This makes SOAR an excellent candidate for simulating cognitive developmental phenomena. Much of the child's procedural knowledge has been described in rule-like terms, and a perennial question has concerned how such knowledge is constructed. I have doubts, however, about whether this initial promise is realized by current SOAR models.

I will touch on a number of general concerns before focusing on developmental models. First, is it appropriate to represent declarative knowledge in condition-action rules? It would seem rather awkward and expensive to do so if the memory system were to be content-addressable. Different rules would be required to encode all the various cues that could potentially access any particular memory. Consequently, one could expect the learning of declarative information to be too complex in such a system, a prediction borne out by the data chunking experiments in SOAR. This is not a model with immediate intuitive psychological appeal.

A second question concerns the admittedly elegant chunking mechanism. Does the chunking of search results at impasses account for all human learning? It is convenient to break this into two subquestions, dealing with the occasions and the mechanisms of learning. although SOAR specifies that learning via chunking occurs only at impasses (i.e., when SOAR lacks knowledge of what to do), it is apparent that humans learn in a wider variety of circumstances. They learn incidentally to problem solving (Nelson 1976) and when problem solving is progressing without impasse (Siegler & Jenkins 1989; VanLehn 1991).

If one wanted to postulate a single basic mechanism for all human learning, one might be better off with association rather than a complex algorithm like chunking that is geared to a particular exploratory mechanism (search) and a particular representational format (condition-action rules). Otherwise, seemingly simple learning problems such as learning declarative information become too convoluted, as they have in SOAR.

A third reservation concerns the rigidly deterministic nature of SOAR, and indeed most other symbolic level models. They do not easily capture the individual differences found in human performance. To account for such variation, one could assume that any part of any rule is a parameter that needs to be fitted to data, but this generates such a huge parameter space for realistic problems that modeling human variation becomes computationally expensive.

Fourth, close inspection reveals that SOAR models typically require a great deal of domain specific knowledge to be built in prior to learning something genuinely new. Theoretically, much of that prior knowledge could be learned through search, but in practice the construction of such prior knowledge often remains an unanalyzed problem. This gives SOAR simulations a highly engineered flavor that to some extent belies the domain generality of the basic ideas in SOAR.

Newell's choice of the balance scale as the first cognitive-developmental domain for SOAR is appropriate. Balance scale experiments, which require the child to predict the result of placing weights at distances to the left and right of a fulcrum, have generated clear, replicable results and a classic progression of stages, each of which is readily describable in rules (Siegler 1981). This has ensured that balance scale regularities now constitute a kind of benchmark for a variety of computational models.

Although the SOAR model of balance scale stages is not fully described in the book, a number of difficulties with this model are apparent. The SOAR model progresses through only the first three of the four stages found with children; it reaches each successive stage far too rapidly compared to children, and the order of its stages may depend entirely on the order of the types of balance scale problems it encounters.

Moreover, like any symbolic rule model, SOAR cannot easily capture the torque difference effect. This is the tendency for problems with large absolute torque differences to be easier for children to solve (Ferretti & Butterfield 1986). (The torque on each side of the fulcrum equals the product of weight x distance on that side). Symbolic rules (e.g., predicting that the side with the larger weight will go down) fail to capture this effect because they typically apply regardless of the problem's torque difference. Rule-based solutions might be contrived, but would probably be awkward. For example, one could construct rules of the form "if torque difference is greater than x, then apply rule i." This would have the system using torque differences to determine which rule to apply well before the system could use torque differences to solve balance scale problems.

In contrast to these limitations, connectionist network models are able to capture all four balance scale stages with plausible variations such as tentative transition points and some degree of stage skipping and regression, as well as the torque difference effect (Shultz & Schmidt 1991). This is not to claim that SOAR models are incapable of simulating balance scale effects and other developmental phenomena, but just to underscore that SOAR models do not currently match the success of some alternative models. The difficulties with the torque difference effect suggest that symbolic level models like SOAR may not provide the best level of description even for relatively high level phenomena such as reasoning about the balance scale.

ACKNOWLEDGMENTS
This work was supported in part by a grant from the Natural Sciences and Engineering Research Council of Canada.

Problem spaces, language and connectionism: Issues for cognition

Patrick Suppes
Department of Philosophy, Stanford University, Stanford, CA 94305
Electronic mail: *suppes@csli.stanford.edu*

Reading Allen Newell's *Unified theories of cognition* has been a great pleasure. It is a book one can argue with and enjoy on almost every page. It has an audacity and sweep that is rare in psychology. Much of what Newell has to say I agree with, but of course it is the function of a commentary of this sort to stress the disagreement. All the same, I shall not be able to refrain from returning at several points to things he has to say that we should all appreciate and learn from. I comment on three general issues.

1. Problem-space hypothesis. Newell emphasizes that the uniform use of problem spaces as the task representation is a central aspect of SOAR, his exemplar of a unified theory of cognition. There is much that is appealing about this, but there is also a serious problem of generality. I am reminded of the set-theoretical semantics of Montague (1974), so popular in linguistics and philosophy a decade or so ago. Without restriction, set theoretical apparatus can always be made rich enough to describe any structure likely to be encountered in language or elsewhere in nature. This is just the problem. It is as if we thought we were making progress in physics by introducing the particle-space hypothesis, that is, the formulation of all physical problems in terms of particles. We have not made significant scientific progress, however, until specific further constraining assumptions have been made and we understand how to apply them in rich detail to many different domains. I am not negative about Newell's introduction of problem spaces as a uniform method of importance (pp. 161 ff.), but I am skeptical as to whether he has at all persuaded us that this in itself has the great significance he claims. An obvious alternative, put forward only briefly in the last chapter, is that human cognition consists of many distinct modules, only a small number of which can be properly regarded as problem spaces. What Newell says about the uniform use of problem spaces is not particularly wrong in all kinds of places, but rather too many of the details are missing. I turn now to one striking example, a strong candidate for modularity.

2. Language learning. In the last chapter a fairly large section is devoted to language. There is an interesting general discus-