Intellectica, 1996/1, 22, pp. 203-227

JALONS

David LEISER* — *Computational emergence and developmental emergence: a sobering survey*[1]

Un robot-poète ne nous fait pas peur. (...) Cultivons, cultivons notre polyvalence. Nous ne sommes pas parfaits, mais très adaptables... Sachons tout... L'avenir est à Pic de la Mirandole. Mirandolez...

Boris Vian (1953)

COMPUTATION AND EMERGENCE

Some fifteen years ago (Leiser,1979), emergence was difficult to conceptualize: developmental psychology consisted of two unconnected components. On the one hand, an increasing number of studies documented the various stages in the acquisition of this or that domain. The description of the stages was couched in natural language, or with the help of algebraic or programming languages. On the other hand, there were some theories, notably Piaget's, that described in general terms the way cognitive development might proceed. It was impossible to say much about emergence, beyond somewhat metaphoric cybernetic statements because there was no way to describe intermediate levels.

Extensive progress in the representation of intermediate levels was made since, and was carried by two main currents. One is modular, « localist ». It analyzes a complete and complex skill as the interplay of a large number of rules (productions, classifiers etc.) that jointly

---

* Department of Behavioral Sciences, Ben-Gurion University, P.O.Box 653, Beer-Sheva, Israel, Email: dleiser@bgumail.bgu.ac.il

[1] The support of the Center for Research in Ergonomics and Safety, and that of the LAFORIA, Université Pierre et Marie Curie, Paris, is gratefully acknowledged. The helpful comments of Daniel Memmi and Anh Nguyen-Xuan significantly improved the presentation of this paper.

generate the required behavior. The emergence of a better adapted behavior may then be described as the gradual assembling and refining of a set of elementary rules, each one of which may, up to a point, be understood in isolation (Anzai and Simon, 1979; Young, 1976). Progress was also made in the study of transition mechanisms, of which a variety have been implemented (e.g., Holland, Holyoak, Nisbett and Thagard, 1986) though it remains unclear how well they can really account for the range of cognitive structures believed to exist (Leiser, 1989).

The second current is formed by the varieties of connectionist models (PDP, auto-association and Boltzman machines, Brain-in-the-Box, Kohonen self-organizing maps, ARC, etc.). These distributed systems have already provided convincing examples of the acquisition of a very diverse array of knowledge.

There remain important difficulties. As has often been noted, the preadaptation of the system's architecture to the domain to be learned vitiates, or seriously relativizes the claim that the system learned on its own. This is especially true when several subsystems are cleverly combined to let a certain behavior emerge, (as in the model by Schyns, 1991, who combined most of the architectures just mentioned to model the emergence of classification and naming behavior). Such systems do indeed demonstrate the emergence of the required *behavior*, but the development of their *architecture* is closer to programming, relying as it does on the combination of existing subroutines or known algorithms to effect modular parts of the task.

Thus, we are in a much better posture than in the past to discuss mechanisms of emergence, inasmuch as we have some understanding of what emerging, incomplete knowledge might look like. At the same time, the proposed systems are yet to demonstrate that they are capable of generating the range of psychological structures that are found to emerge. This will be especially difficult to the extent that the grain of the elements relied on by the mechanisms is small, the scope of what those systems are trying to learn is large, and there are fewer inbuilt assumptions about the domain to be acquired.

While the progress on these fronts is beyond dispute, I propose to discuss its limitations by sketching the magnitude of the task facing developmentalists. When such a survey is made, the limits of what has been achieved become painfully apparent. The extent of emergence and our success in modeling can be decomposed into three logically ordered questions: *conceptual*: Where should we look for knowledge? e*mpirical:* Is it in fact found in human development (e.g., do Piaget’s

overarching operational structures really exist)? and *computational*: How can we model that emergence? The present survey will mainly be concerned with the first of these levels (the « knowledge level », that of Piaget's epistemic subject) because this level is what sets apart a theory of emergence.

The importance of the knowledge level is indeed widely acknowledged, even by that most pragmatic of Artificial Intelligence research domains, knowledge engineering (Schreiber, Wieling and Breuker, 1993). However, how autonomous is it ? Can behavior be usefully studied without detailed anatomic or computational analysis ? Can one, for instance, usefully separate a distinct knowledge level in the case of a neural network that performs a task successfully ? Coltheart (1994) argues that one can and indeed must. To him, the exemplary conflict between his own psychological research (on the « dual route » model) and recent successes in modeling the same phenomenon by means of neural networks, is to be resolved by contrasting network architecture and functional architecture, and not letting the former distract one from studying the latter :

> « The network architecture is determined by the modellers, and is obvious. The functional architecture (...) can be very difficult to discover in realistically large networks. What are cognitive psychologists interested in ? Well, obviously, functional architecture. » (p.21)

Others hold a view closer to the one presented here, and see the two architectures as complementary (Bates and Elman, 1992; Leiser, 1979 ; Plunkett and Sinha, 1992; Clark, 1992). To ground the attribution of a structure to a subject, one must show how development necessarily causes the structure to emerge, and conversely, why the structure is an attractor for the developmental process, as expressed in Piaget's dictum *Pas de structure sans genèse, pas de genèse sans structure* (no structure without development, no development without structure). A complete explanation would therefore involve all three terms : the emergent structure, the architecture of the learning mechanism, and a chronicle of the development. The structure sets the goal, with reference to the system's internal or external adaptation. The architecture accounts for how the system comes to construct the structure, and thereby explains the chronicle. A very similar view was put forward by McClelland and Seidenberg (1992 ; see also Plaut, McClelland, Seidenberg and Patterson, 1994) who describe their approach as follows : they postulate a specific way of representing the stimulus world, a learning rule, and an architecture, then observe that under those conditions, a certain type of knowledge representation

necessarily results. It is not built in, but arises from interactions among independent factors and this, to the extent it succeeds, provides the explanation for development. The study of emergent developmental properties and of the course of this emergence, precede the modeling work, whose task is to show why certain structures are acquired.

On these points, theorists like Piaget had little more to say than that « sooner or later » the next stage would necessarily be reached. Piaget's account was explicitly at the level of his «functional invariants» as opposed to that of specific mechanisms (Piaget, 1975, pp. 180-181). The distinction is indeed valid, and the possibility that a variety of mechanisms may fulfill the same function is well known. But as long as *no* mechanism is forthcoming, the account lacks a key component. The newer computational models strive to correct this deficiency. With this in mind, it seems useful to list the range of phenomena that have, ideally, to be accounted for, though not necessarily by independent mechanisms.

LOCI OF EMERGENCE

The search for such knowledge structures involves scouring behavior to identify the loci where knowledge is found or shaped, more often than not in implicit form. Putative cognitive structures have very different scopes, and the way to conceive of them, and hence to model them is correspondingly different. Let us arbitrarily separate them into small- and large scope structures.

For the smaller scope, knowledge is implicit in data and associated procedures, in the representational formats, and in the system's architecture. One important part of development could be the transfer of knowledge between these levels (Leiser, 1987; Moshman, 1990; Clark, 1992; Karmiloff-Smith, 1993). For the larger scope, knowledge consists in the adequation and organization of categories of thought (Kant, Piaget), ontogenies (Keil 1979), inborn naive theories (Carey 1975) and the ways and conditions for differentiating or extending existing ones.

In a case study of seriation procedures (Leiser and Gillièron, 1990) we showed that the loci of logic are multiple and heterogeneous. Logic is found in rules for symbolic transformations, in the overall coherence of the procedure, in the ability to assimilate a given problem to an appropriate approach, and in the accommodatory potential of the procedure. In more detail:

Symbolic representation: Some simple inferences are readily described in terms of aspects of mental manipulation of symbolic representations. Thus, elementary reasoning steps may be directly based on basic logical rules. This is the oldest approach to the simulation of reasoning, going back to the notion of uniform resolution algorithms. The elementary rules accounting for the transformational steps may also be more specific to the domain at hand: «General methods are weak methods ».

Overall coherence: The overall coherence of a procedure embodies another aspect of operative knowledge. While it is true that elementary production rules, say, may have their own, modular function, they co-evolved jointly under adaptive pressure, and their adaptation to a given task is a joint property. Moreover, the entire adaptive package may be transferred to an isomorphic domain, showing that its coherence may properly be ascribed to the system.

This is very clear in Genetic Algorithms (Holland *et al*., 1986; Koza, 1991; Belew *et al*. 1991; Boers *et al*., 1993), by far the most successful model for the unguided development of production systems. Reinforcement of successful rules is calculated to reinforce successful cooperation between existing rules, rather than merely successful individual ones (as in Anderson, 1986; Klahr, Langley and Neeches, 1987).

Assimilation: Assimilation is a basic form of understanding. The use of a given representational format, procedure or architecture betrays or manifest an implicit belief in its appropriateness. The contents of this belief can be made explicit, and this is routinely done in software engineering, whenever an orderly attempt is made to define generic problem-solving methods, with an eye on reusing them on a wider range of problems (David, Krivine and Simmons, 1993). An essential aspect of development may be the exploitation of regularities in the data encoded at the symbolic level, to identify existing patterns, or even to construct representational formats that will embody those regularities in their very architecture . Such a passage could involve the extraction of the regularities as the basis for building a new representational architecture, by some identifiable process of reflective abstraction or rely on a distributed representation, where the distinction between the representation of rules and that of exemplars becomes moot (Shultz, Schmidt, Buckingham, & Mareshal, 1993).

It may be useful to recall the discussion levels we introduced before: conceptual, empirical and computational. We are considering the relations between regularity patterns in the structure of already encoded

data, and the subsequent ease of encoding of additional data obeying the same structural relationships, as a conceptually identifiable locus of emerging knowledge. If learning novel cases similar to previous ones is shown to be facilitated by knowledge of the previous ones, this is an important psychological principle. Finally, how this functional regularity is implemented is the computational question.

Potential for Accommodation: The last and most subtle of the loci is the accommodation potential of the assimilatory schema, representational format, or architecture. We noted that the use of any of these structures implies an implicit belief in its appropriateness, hence a specification of the range of application of the structure. However, an assimilatory structure may evolve, under influence of the occasions on which it was used. There is no telling which way it will drift, since this will depend on the history of its encounters with related problems. But it will not drift without restraint, and its evolution must accord with more abstract structural constraints.

THE SCOPE OF EMERGENCE

A common observation about AI is that its successes concern toy problems, « micro worlds ». Most cognition involves much larger systems (Guha and Lenat, 1990), hence these practical successes are of debatable general significance (Schank, 1991). The same may be true of psychology. The proper scope of a cognitive domain remains very unclear. How widespread is the influence of the structure of one domain? Is it limited to that domain only? To isomorphic or analogical ones? To ones that are associated by their content or context? To ones belonging to the same « frame of mind » (Gardner, 1983)? Very little is known about these questions. In what follows, I discuss some possible scopes, from smaller to larger. Additional, cross-cutting classifications are certainly possible.

Individual concepts and relations

The smallest scope is of course that of individual concept formation. This level was much studied both in psychology and in Machine Learning (Dietterich, 1986; Estes, 1986) At the same level belongs the extraction of individual patterns of co-occurrence. However, even simply determining frequencies or patterns of co-occurrences ("intuitive covariation", Wright and Murphy, 1984; Jones and Smith, 1993) is influenced by existing theories and expectations. "Pure" cases are rare, and perhaps atypical as well.

Categorization

Here, the goal is to classify correctly several concepts in contrast to one other. This involves two aspects: categorization (creating novel categories) and classification of new items in those categories. The domain has been extensively studied in classic AI, Machine Learning and in connectionism, and a range of models are now available (Medin, Wattenmaker, & Michalski, 1987; Harnad, 1987).

While the topic has also been extensively studied by psychologists, both in children and in adults, remarkably little firm knowledge is available (Ahn & Medin 1992; Ward Vela & Hass, 1990; Murphy and Medin, 1985). In particular, the choice between classic, exemplar, prototype, family resemblance and distributed models is very much an embarrassment of riches. Faced with the straightforward question: « how does the child learn to categorize its experience? » developmentalist cannot say much that is definite.

Organization in domains

Domain types

Beyond individual concepts and mutually contrasting categories, matters become even less definite. What is a domain? Surveying the literature, it is obvious that there are many types, and as many approaches to defining them. A rapid sample: logico-mathematical structures (Piaget, passim), domain-bound intelligences (Gardner, 1983), emerging frame systems (Schank, 1982), coherent theories (Carey, 1991; Forbus and Gentner, 1986; Keil, 1992; Springer, 1990); systems of mutual definitions (Piaget, Vygotsky), cooperating rule clusters (Holland, Holyoak, Nisbett, & Thagard, 1986; Anderson, 1983; Laird, Newell, & Rosenbloom, 1987; Klahr, Langley, & Neeches, 1987).

Different domains are carved out by the diverse techniques used to study them: memory experiments, reasoning, priming, logical coherence, peaks in development, psychometrics, and neurologic data ... The ability of theorists to define a domain by reference to one of those sources, then discuss its development in relative isolation should give us pause: there is a danger of confusing natural kinds with domains delineated for methodological or theoretical reasons. The basic conceptual question (*do we know what types of domains to look for?*) is far from resolved. The following tripartite classification may be useful:

1. *Imposed domains* have their coherence defined relative to some theoretical considerations of method -- i.e., they are domains for the psychologist or the theoretician, not for the system studied.
2. *Systemic domains*. The organism defines cognitive domains, due perhaps to innate predispositions (e.g. neuronal and/or computational, such as stereoscopic *v*. kinematic sources of 3D information).
3. *Pragmatic domains*: the pattern of mutually relevant knowledge may not be restricted to a single systemic domain. Fodor (1983) described the highest knowledge level as one where everything may, in principle, be related to everything else. But of course, pieces of knowledge are not equally relevant to one another, and relative modularity may emerge from patterns of co-activation.

Design Requirements

Even when we focus on substantive domains, there remain major conceptual and empirical questions. In the evolution of a domain, several design requirements must be balanced: adaptation, continuity, and (re)structuring. The empirical question is then what balance is in fact struck in a given domain, by a given individual, at a given time, and what are the mechanisms responsible for reaching this balance.

1. *Adaptation*. The domain should relate to the environment appropriately. It should capture regularities in the environment and lead to adapted behavior. This is of course the basic requirement, that defines development as such, as opposed to mere drift.

2. *Continuity with what exists*. New knowledge merges with what is already known. "Gracefulness" is essential for an organism whose exchanges with the environment must be maintained. (cf. Piaget's fundamental equilibration cycle, 1971; Gold, 1987; Jacob, 1970).

3. *Structuring*. We just saw that new knowledge should take into account the implicit structure of existing knowledge. But conversely, accumulated knowledge may be exploited to improve existing representations. There are several ways to do this.

   - *Streamlining*: the same representational format is retained, but the information is reorganized more efficiently. Generalizations may be formed, exceptions identified, complicated formulations simplified. This is restructuring in the weak sense, as discussed by Carey (1991).

   - *(Re)structuring:* This too is a transition to a better representation, so that the same set of instances, behavior, classification, may be generated more efficiently, but the representational format itself is replaced or modified, as better advantage is taken of regularities and lawfulness in the data being represented.

     The tradeoff is between the short-term benefit of getting something to work correctly in the environment, be it by a succession of patches, and the longer term advantages of a more efficient, more predictive and more stable representation. This possibility is rarely discussed in developmental work, but was seen by Piaget (1974, on « cognitive phenocopy »; Leiser, 1987) as a central feature of development.

   - *Grounding:* Knowledge implicit in the functioning of a system (whether regularities in the data represented or the organization of its representation) may come to be extracted and explicitly represented in some meta-language, and this presents various

advantages: that knowledge becomes more open to scrutiny, amenable to various transformations (logic, language) and may become organized in a principled grounding of the behavior, as complete theories are developed, which in turns may guide the orderly growth of the domain. Over the years since Dreyfus (1986) maintained that skills are not representable by rules at all ("experts know the spirit of the rules and therefore know when to break them"), his point has gained acceptance. It is now clear that expert systems must be endowed with proper grounding of their rules (Forbus and Gentner, 1986; David, Krivine and Simmons, 1993). Pure rule-based systems cannot represent sufficiently the knowledge domain, and are therefore unlikely to be successful if they are allowed to grow or modify themselves without guidance from some other mechanism -- whether a teacher or an internal, conceptual model. This transition is usually discussed under the head of the transition from implicit to explicit representation.

The three requirements function dialectically, and there is a tension between continuity and adaptation, on the one hand, and the need for structuring, on the other. A system that develops using the patching approach may be effective, up to a point, but this strategy may result in cognitive impenetrability: one reason why sub-systems are modular and encapsulated may be that the other systems cannot figure them out! Whereas knowledge compilation goes from declarative to procedural (e.g., Anderson, 1983), an important counter-direction goes from procedural to declarative, with increased control (Tzelgov, Henik, & Leiser, 1990; Karmiloff-Smith, 1993; Clark, 1992).

Failure to restructure may have its costs. Spurious pseudo-information or pseudo-rules that never served any function, but happened to be associated in the right way with the efficient causes, may be retained. Outdated rules that did serve some purpose in the past may be retained even when the original reasons for including them are no longer applicable. This is nicely illustrated by Primo Levi's wistful musing:

> « ... but formulas are as holy as prayers, decree laws and dead languages, and not an iota in them can be changed. And so my ammonium chloride (...) by now completely useless and probably a bit harmful, is religiously ground into the chromate anti-rust paint on the shore of that lake, and nobody knows why anymore ». (Levi, 1984, p. 159).

In the absence of principled reasons for including or removing rules, pointless and sometimes harmful rules will be retained. But structuring presents hazards of its own: once a representation or an architecture

embodies some assumption about the material being represented, that assumption becomes entrenched, and hard to change or even question. Grounding too has its drawbacks: to the extent that knowledge was incompletely or erroneously extracted, performance deteriorates. U-shaped development (Strauss and Stavy, 1982) suggests that a change of structuration is taking place.

These various restructuring possibilities may be more tightly related than is commonly realized. Extracting regularities, devising better representations and constructing better theories are all interrelated, and bound with the notions we entertain about the range of possible mental representation and their ontological and methodological status. Suffices it to mention the revolution brought about by the development of distributed representations.

There is very little empirical data on any of these issues. Does restructuring ever take place? And if so, is it of the strong or of the weak kind, that is, is it the equivalent of a paradigm shift, or is it limited to the enrichment or reorganization of the existing data base? (Carey, 1991; Keil, 1986).

Beyond single domains

Domains do not evolve in isolation. Learning a $n+1^{th}$ domain is different from learning the first n. And knowing several domains has its impact on the individual domain. A universal novice is a novice in a way a domain-specific novice is not, for learning previous domains may have had its effect.

Known domains serve as the source of analogies, as templates in the construction of the new domain, once both are sufficiently structured to detect the relevance of an analogy (Goswami, 1991; Holyoak, Novick & Meltz, 1994; Gentner and Toupin, 1986; Gentner, Ratterman, & Forbus,1993; Brown, 1989). A domain or scheme might be more or less profoundly affected by a successful analogy to which it gave rise. In a way, the most difficult issue is to understand how these degrees might merge into one another. This is of course yet a different form of the emergence issue we mentioned at the onset.

1. *Transfer* is the weakest linkage. It need have no lasting effect. The original scheme remains as it was, and the cognitive structure, the scheme, is not lastingly modified by the encounter. No variabilization or abstraction takes place. There is only the ephemeral act of setting up the analogy, and possibly the updating of an associative link, that will make the analogy more likely in the future.

Even this weak linkage raises difficult indexation problems, too easily forgotten except in software engineering research: finding a relevant, potentially useful matching supposes some indexation scheme, some ability to initiate a search for relevantly similar items. This is no trivial task (Medin, Goldstone & Gentner, 1993). Leiser and Gillièron (1990) analyzed the significant degree of understanding required before transfer may take place, whereas others showed that transfer may be initiated more readily on the basis of superficial resemblance than similarity of structure (Reeves and Weisbert, 1994, Novick and Hmelo, 1994; Gentner and Youpin, 1986; Gentner, 1989). Bassok and Holyoak (1989) for instance, showed that some interdomain transfers of structure are unidirectional. Whereas the questions are now more sharply defined than in past, little remains known about transfer and analogy.

There are very few ideas about mechanisms that might effect this, but there are some beginnings, even when the structure is as diffuse as in a back-propagation network (e.g., Pratt, 1993).

2. *Schema abstraction by progressive variabilization*. Transfer is only within a zone of mental elasticity, the assimilatory region of the scheme (Vygostky, 1934/1986). The zone itself is fuzzy: increased effort manages to effect the transfer to more distant items. The implicit definition of what the scheme does evolves as that zone extends. The original paradigmatic case need not remain central or retain a special status, as it becomes subsumed under the more abstract scheme that evolves. When this process has succeeded far enough, one talks more naturally of abstraction.

   Reasoning relies on tools that are partly content-bound, and partly abstract (Reuchlin, 1973). Even in adults, freedom from context and content is far from complete. Reasoning schemata are not entirely contentless, and there are pragmatic reasoning schematas of various descriptions (Braine, 1990; Cheng and Holyoak, 1985).

3. *Complete abstraction*. The ultimate end-result of this process is when structurally equivalent problems are solved with equal ease. No need to recall A in order to solve B in the same way, once the common structure has been abstracted. This process is gradual, not a single event. It occurs progressively, at the occasion of interactions with the environment. At this level, if it is found, there is no longer any *decalage*, and information may be written generically in some metalanguage. But, does this really exist? Paradoxically, it is easier to devise an abstraction mechanism than one whose product remains

partly content-bound. Do Piagetian operational structures have psychological validity, or are they theoretical idealizations only?

4. *Coordination of domains* The same empirical situation may give rise to contradictions, to incoherence, if assimilated to different domains simultaneously. Conflict resolutions go from *demarcation* to complete *merging* with attendant conceptual change. The former leaves the two domains as they were, merely adding with the ability to discriminate when either of them is appropriate. The latter may involve a complete reworking of the two domains involved. Piaget (1975) viewed the need to coordinate one's mental structures as one of the two great sources of development (the other being the need to adapt to the external environment).

Functional domains or meta-skills

Beyond the content, as the universal novice becomes a seasoned cognizer, his/her thought and developmental processes may develop in a way that is not tied to a specific domain. Learning N domains may develop heuristic skills. The work in Machine Learning is very useful to develop an appreciation for the range of meta-skills that are in principle applicable to the issue of learning and development. Development may take place at this level, regardless of whether these should be considered as implemented procedures or as abstract functions, whose implementation is diffuse in many processes. The skills can be used initially to organize a new substantive domain, but initially only: general methods of formal induction are notoriously very inefficient; eventually, every domain should develop its own heuristics, since domains are often too different from one another. The "engines" of expert systems are good examples: checking hypotheses, differential diagnosis, backward chaining, and forward chaining, probability combinations, backtracking, constraint satisfaction, minmax, blackboard strategies etc. are all examples of relatively general skills. To do justice to this topic, a survey equivalent to that on substantive domains would be required. I will however restrict myself to two examples of possible "functional domains".

1. *Analogy drawing itself* may be such a domain-general skill. Does one become more proficient at drawing analogies with practice? Does one become better at it, using for example better framed retrieval cues? There is very little evidence for this. Reasoning was shown to be generally more mature in any number of ways when the domain was well-known to children (Chi and Ceci, 1987; Chi,

Hutchison, & Robin, 1989) and *decalage* in every cognitive domain appears to be the rule.

2. *The development of argumentative and reasoning skills.* Vigotsky (1934/1986) saw reliance upon definitions, conscious deliberate use of concepts in a system of contrasting concepts as a general skill that develops and is widely applicable. The known effect of schooling on cognitive style suggests he may have been right. In the same vein, Moshman (1990) distinguishes four staggered stages where content, inference, logic and metalogic all progress from implicit to explicit. But again, depressingly little is known.

Rosser (1994) distinguishes three developmental strands with different onset time and developmental pathways:

ï *domain-specific knowledge acquisition:* of rudimentary format and innately specified, it forms the core.

ï *expertise specific knowledge acquisition,* that must follow domain-specific knowledge acquisition.

ï *domain-general cognitive operations*: facts, strategies, procedures, operations, algorithms, reasoning schemes, these are neither content nor context bound. Where this originates is not clear, nor do we know its limits, but this is the strata that Piaget sought to explain.

The first two of those "strands" recall our discussion of design requirements for the learning of domains, while the last refers to knowledge that cuts across domains. Following her survey of the literature on the third, Rosser was lead to conclude:

> « Where does explicit reasoning or metalogic come from? What are the developmental mechanisms that enable access to it. ... We know deductive reasoning is age related, task related, content related and knowledge related. Whether, if at all, a domain-general cognitive function is implicated still remains unclear. »

THE CHALLENGE AHEAD

There has been considerable progress in the understanding of emergence, starting from the information-processing revolution, through the more recent neural networks and other neuronally-inspired conceptions (such as Abeles, 1990; Shastri & Ajjanngadde, 1993). However, this progress should not obscure the very narrow scope of existing models. The alternatives are at present models of the emergence of small-scale abilities (the brain as tool-box, as Robin

Campbell put it); or the still largely programmatic notion of a multitude of minute modular items that combine to form higher, more complex structures (Holland Holyoak, Nisbett, & Thagard, 1986).

On the empirical side, not enough is known about such emergence. We mentioned a wide range of cognitive loci where emergence may be required. By and large, it is not known whether it actually takes place. Several particularly difficult issues were raised, and situated in relation to the entire domain of cognitive development. In particular, we asked :

- Does restructuring take place?
- How do large-scale domains, schemes and skills form?
- How do domains and schemes coordinate with one another?

These issues are profoundly interrelated. Making progress -- conceptual, empirical and computational -- on these questions is the true challenge of developmental theory.

Scientific reading suffers from an obvious sampling bias. Reports concern achievements: data analyzed, programs constructed. One may therefore be forgiven for forming the impression that cognitive emergence is more or less understood. It is only when a deliberate effort is made to map what should be known that the true state of affairs emerges: as this survey strove to show, the challenge of developmental theory is still ahead of us.

David LEISER

Bibliography

Abeles, M. (1991). *Corticonics*. Cambridge, MA: Cambridge University Press.

Ahn, and Medin, D.K. (1992). A two-stage model of category construction. *Cognitive Science, 16(1)*, 81-121.

Anderson, J. R. (1983). *The Architecture of Cognition*. New York: Freeman.

Anderson, J. R. (1991). The adaptive nature of human categorization. *Psychological Review, 98*, 404-429.

Anderson, J.R. (1986). Knowledge compilation: the general learning mechanism In Michalski R.S., J.G. Carbonnell and T.M. Mitchell (Eds.) *Machine Learning: An Artificial Intelligence Approach* vol. 2. Palo Alto: Tioga Press.

Anzai, Y. and Simon, H.A. (1979). The theory of learning by doing. *Psychological Review, 86,* 124-140.

Bakker, P. (1994). Good models of humble origins. In J. Wiles, C. Latimer, and C. Stevens (Eds.), *Conference on Connectionist modelling and psychology*, Brisbane: Department of Psychology, University of Queensland.

Bassok, M., and Holyoak, K. J. (1989). Interdomain transfer between isomorphic topics in algebra and physics. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 15(1)*, 153-166.

Bates, E. A., and Elman, J. L. (1992). Connectionism and the study of change. In M.H. Johnson (Ed.), *Brain Development and Cognition: A reader*. Oxford: Blackwell.

Belew, R. K., McInerney, J., and Schraudolph, N. N. (1991). Evolving Networks: Using the Genetic Algorithm with Connectionist Learning. In C.G. Langton, C.Taylor, J. D. Farmer, and S. Rasmussen (Ed.), *Artificial Life II*(Vol. X,). New York: Addison-Wesley.

Boers, E.J.W., Kuiper, H., Happel, B.L.M., and Sprinkjuizen-Kuyper, I.G. (1993). *Designing modular artificial neural networks* (No. 93-24). Department of Computer Science Leiden University.

Braine, M.D.S. (1990). The "natural logic" approach to reasoning. In W.F. Overton (Ed.), *Reasoning, logic, and necessity* (pp. 135-158). Hillsdale, NJ: Erlbaum.

Brown, A. L. (1989). Analogical learning and transfer: What develops? In Vosniadou, S. and A. Ortony (1989). *Similarity and analogical reasoning*. (pp. 369-412).Cambridge, MA, Cambridge University Press.

Carey, S. (1985). *Conceptual change in childhood*. Cambridge, MA: MIT Press.

Carey, S., and Gelman, R. (Ed.). (1991). *Epigenesis of Mind: Essays in biology and knowledge*. Hillsdale NJ: Erlbaum.

Ceci, S.J. (1989). On domain-specificity -- More or less general and specific constraints on cognitive development. *Merrill-Palmer Quarterly, 35*, 131-142.

Cheng, P.W., and Holyoak, K.J. (1985). Pragmatic reasoning schemas. *Cognitive Psychology, 17*, 391-416.

Chi, M., and Ceci, S.J. (1987). Content knowledge: its role, representation and restructuring in memory development . In Reese, H.W. (Ed.), *Advances in Child Development and Behavior* (Vol. 20): Academic Press.

Chi, M.T.H., De Leeuw, N., Chiv, M. H., and LaVancher, C. (1994). Eliciting self explanations improves understanding. *Cognitive Science, 18*, 387-439.

Chi, M.T.H., Hutchison, J. E., and Robin, A. F. (1989). How inferences about novel domain-related concepts can be constrained by structural knowledge. *Merrill-Palmer Quarterly, 35*, 27-62.

Clark, A. (1990). Connectionism, competence and explanation. *British Journal for the Philosophy of Science, 41*, 195-22.

Coltheart, M. (1994). Connectionist modelling and cognitive psychology. In J. Wiles, C. Latimer, and C. Stevens (Ed.), *Proceedings of Conference on Connectionist modelling and psychology*, Brisbane: Department of Psychology, University of Queensland.

David, J.-M., Krivine, J.-P. and Simmons, R. (1993) (Eds.): *Second Generation Expert Systems*, Berlin: Springer.

Dietterich, T. G. (1986). Learning at the knowledge level. *Machine Learning, 1*(3), 287-316.

Estes, W. K. (1986). Array models for category learning. *Cognitive Psychology, 18*, 500-549.

Fodor, J. A. (1983). *The modularity of mind*. Cambridge, MA: MIT Press.

Forbus, K.D. and Gentner, D. (1986) Learning physical domains -- toward a theoretical framework. In R.S. Michalski, J.G. Carbonell and T.M. Mitchell (Eds.) *Machine Learning: An Artificial Intelligence Approach , vol 2*. Palo Alto: Tioga Press.

Gardner, H. (1983). *Frames of mind: the theory of multiple intelligences*. New York: Basic.

Gentner, D. (1989). The mechanisms of analogical learning. In Vosniadou, S. and A. Ortony (1989). *Similarity and analogical reasoning*. (pp. 199-241). New York, Cambridge University Press.

Gentner, D. and C. Toupin (1986). "Systematicity and surface similarity in the development of analogy." *Cognitive Science, 10*, 277-300.

Gentner, D., Ratterman, M. J., & Forbus, K. D. (1993). The roles of similarity in transfer: separating retrievability from inferential soundness. Cognitive Psychology, 12, 306-355.

Gold, R. (1987). *The Description of Cognitive Development -- Three Piagetians Themes*. Oxford: Oxford University Press.

Goswami, U. (1991). Analogical reasoning: what develops? A review of research and theory. Child Development, 62, 1-22.

Guha, R. and Lenat, D. B. (1990). CYC: a midterm report. *Ai Magazine*, 32-59.

Harnad, S. (1987). Category induction and representation. In Harnad. S. (Ed.), *Categorical Perception*. Cambridge, UK: Cambridge University Press.

Holland, J. H., Holyoak, K. J., Nisbett, R. E., and Thagard, P. R. (1986). *Induction - Processes of Inference, Learning, and Discovery*. Cambridge, MA: MIT Press.

Holyoak, K. J., and Thagard, P. (1989). Analogical mapping by constraint satisfaction. *Cognitive Science, 13*, 295-355.

Holyoak, K. J., Novick, L. R., & Meltz, E. (1994). Component processes in analogical transfer: mapping, pattern completion and adaptation. In K. J. Holyoak & J. A. Barnden (Ed.), Advances in connectionist and neural computation theory.(Vol. 2: analogical connections, pp. 113-180). Norwood, NJ: Ablex.

Jacob, F. (1970) *Le Jeu des Possibles* Paris: Gallimard.

Jones, A. and Smith, L.B. (1993). The place of perception in children's concepts, *Cognitive Development*, *8(2)* 113-139.

Karmiloff-Smith, A. (1992). Nature, Nurture and *PDP*: Preposterous Developmental Postulates? *Connection Science*, 4(3), 253-269.

Karmiloff-Smith, A. (1993). *Beyond modularity: a developmental perspective on cognitive science*. Cambridge, MA: MIT Press.

Keil, F. C. (1979). *Semantic and Conceptual Development: an Ontological Perspective*, Cambridge, MA: Cambridge University Press

Keil, F. C. (1986). The nonrepresentative nature of representational change: Some possible morals to draw from Nelson's « Making Sense ». *Cognitive Development*, *1*, 281-291.

Keil, F. C. (1990). Constraints on constraints: surveying the epigenetic landscape. *Cognitive Science*, *14*, 135-168.

Keil, F.C. (1991) The emergence of theoretical beliefs as constraints on concepts. In Carey, S., and Gelman, R. (Eds.) *Epigenesis of Mind: Essays in biology and knowledge*. Hillsdale NJ: Erlbaum.

Klahr, D. Langley, P. and Neeches, R. (Eds.) (1987) *Production system models of learning and development*. Cambridge, MA: MIT Press.

Koza, J.R. (1991) *The genetic programming paradigm: genetically breeding populations of computer programs to solve problems.* Computer Science Department, Stanford University.

Laird, J. E, Newell, A. and Rosenbloom, P. S. (1987). SOAR: an architecture for general intelligence. *Artificial Intelligence*, *33*, 1-64.

Leiser, D. (1979) The function of structural representations. In B. De Gelder and P. Geert (Eds.) *Proceedings of the NIAS Conference on Knowledge and Representation*. Wassenaar, The Netherlands.

Leiser, D. (1987). The changing relations of representation and cognitive structure during the development of a cognitive map. *New Ideas in Psychology*, *5*, 95-11.

Leiser, D. (1989) A review of « Induction: Processes of Inference, Learning and Discovery » by Holland *et al.*, 1986. *Journal of Mind and Behavior*, *10*, 303-306

Leiser, D. (1990) Evolution, development and learning cognitive science. *Behavioral and Brain Sciences*, *13(1)*, 80-81.

Leiser, D. and Zilbershatz, A. (1989). The Traveler - a computational model of spatial network learning. *Environment and behavior*, *21*(4), 435-463.

Leiser, D., and Gillièron, G. (1990). *Cognitive science and genetic epistemology - a case study of understanding*. New York: Plenum.

Levi, P.(1984). *The Periodic Table*. New York: Schocken.

Lewis, C. (1987). Composition of productions. In Klahr, D. Langley, P. and Neeches, R (Eds.), *Production system models of learning and development*. Cambridge, MA: MIT Press.

Medin, D. K., Goldstone, R. L., and Gentner, D. (1993). Respects for similarity. *Psychological Review, 100(2)*, 254-278.

Medin, D.K. Wattenmaker, L. and Michalski, R. (1987). Constraints and preferences in inductive learning: An experimental study of human and machine performance. *Cognitive Science, 11(3)*, 299-339.

Moshman, D. (1990) The development of metalogical understanding. In W.F. Overton (Ed.) *Reasoning, necessity and logic: developmental perspectives.* Hillsdale, NJ: Erlbaum

Murphy, G. L., and Medin, D. K. (1985). The roles of theories in conceptual coherence. *Psychological Review, 92*, 289-316.

Novick, L. R., & Hmelo, C. E. (1994). Transferring symbolic representations across nonisomorphic problems. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 20, 1296-1231.

Piaget, J. (1971) *Biology and Knowledge.* Chicago: Chicago University Press

Piaget, J. (1974). *Adaptation vitale et psychologie de l'intelligence: Sélection organique et phénocopie*. Paris: Hermann.

Piaget, J. (1975). *L'Equilibration des structures cognitives: Problème central du développement. Etudes d'Epistémologie Génétique*, vol 33. Paris: Presses Universitaires de France.

Plaut, D. C., McClelland, J. L., Seidenberg, M. S., and Patterson, K. E. (1994). *Understanding Normal and Impaired Word Reading: Computational Principles in Quasi-Regular Domains* (Technical Report No. PDP.CNS.94.5). Carnegie Mellon University.

Plunkett, K., and Sinha, C. (1992). Connectionism and developmental theory. *British Journal of Developmental Psychology, 10*, 209-254.

Pratt, L. Y. (1993). *Transferring previously learned back-propagation neural networks to new learning tasks.* Unpublished Doctoral Dissertation, Rutgers - The State University of New York.

Reeves, L. M., & Weisbert, R. W. (1994). The role of content and abstract information in analogical transfer. *Psychological Bulletin*, 115, 381-400.

Reuchlin, M. (1973). Formalisation et réalisation dans la pensée naturelle: une hypothèse. *Journal de Psychologie Normale et Pathologique, 70*, 389-408.

Rosser, R. (1994). *Cognitive development - Psychological and biological perspectives*. Boston: Allyn and Bacon.

Schank, R. (1982) *Dynamic Memory - a theory of reminding and learning in computers and people*. Cambridge,UK: Cambridge University Press

Schank, R. C. (1991). Where's the AI? *AI Magazine*, 38-49.

Schreiber, G. Wielinga B., and Breuker, J. (1993) (eds.): *KADS. A Principled Approach to Knowledge-Based System Development*, *Knowledge-Based Systems*, vol 11. London: Academic Press.

Schyns, P. G. (1991). A modular neural network model of concept acquisition. *Cognitive Science*, *15*, 461-508.

Seidenberg, M. S., and McClelland, J. L. (1992). *Connectionist Models of Explanatory Theories in Cognition* (Report No. PDP.CNS.92.4). Department of Psychology, Carnegie Mellon University. Pittsburgh.

Shastri, L., and Ajjanagadde, V. (1994) From simple associations to systematic reasoning. *Behavioral and Brain Sciences*.

Shultz, T. R., Schmidt, W. C., Buckingham, D., and Mareshal, D. (1994). Modeling cognitive development with a generative connectionist algorithm. In T. Simon and G. Halford (Ed.), *Developing cognitive competence: New approaches to process modeling*. Hillsdale, NJ: Erlbaum.

Springer, K. (1990). In defense of theories. *Cognition*, *35*, 293-298.

Strauss, D., and Stavy, R. (1982) (Eds.) *U-Shaped Behavioral Growth*. New York: Academic.

Tzelgov, J. Henik, A. and Leiser, D. (1990). Automaticity and controllability in a Stroop task with bilinguals. *Journal of Experimental Psychology: Language Memory and Cognition*, *16*, 760-771.

Vian, B. (1980/1953). Un Robot-poète ne nous fait pas peur. In Vian, B. *Je voudrais pas crever*. Paris: Pauvert.

Vigotsky, L. S. (1934/1986). *Thought and Language* (A. Kozulin, Trans.). Cambridge, MA: The MIT Press.

Ward, Vela, and Hass (1990). Children and adults learn family resemblance categories analytically. *Child Development*, *60*, 214-224

Wright, J. C., and Murphy, G. L. (1984). the utility of theories in intuitive statistics. The robustness of theory-based judgments. *Journal of Experimental Psychology General*, *113*, 310-322.

Young, R. M. (1976). *Children's seriation behavior - an Artificial Intelligence analysis of a Piagetian task*. Basel: Birkhäuser .

***