SIMULATIONS OF STAGEWISE DEVELOPMENT WITH A SYMBOLIC ARCHITECTURE

FERNAND GOBET

To appear in Tschacher, Wolfgang and Dauwalder, Jean-Pierre (Eds.). Dynamics, Synergetics, Autonomous Agents - Nonlinear Systems Approaches to Cognitive Psychology and Cognitive Science. Singapore: World Scientific.

SIMULATIONS OF STAGEWISE DEVELOPMENT WITH A SYMBOLIC ARCHITECTURE

FERNAND GOBET

ESRC Centre for Research in Development, Instruction and Training

and Artificial Intelligence Research Group

Department of Psychology

University of Nottingham

Nottingham NG7 2RD England

frg@psyc.nott.ac.uk

Abstract

This chapter compares Piaget's theory of development with Feigenbaum and Simon's (1962, 1984) EPAM theory. An attempt is made to map the concepts of assimilation and accommodation in Piaget's theory onto the concepts of familiarisation and accommodation in EPAM. An EPAM-like model of the balance scale task is then presented, with a discussion of preliminary results showing how it accounts for children's discontinuous, stage-like development. The analysis focuses on the transition between rules, using catastrophe flags (Gilmore, 1981) as criteria. It is argued that some symbolic models may be described as dynamical systems, in the same way as some non-symbolic models.

1 Introduction

Although this is still an area of controversy, recent research supports Jean Piaget's suggestion that cognitive development occurs in a non-linear, stagewise fashion (Raijmakers, van Koten & Molenaar, 1996; Thomas & Lohaus, 1993; however, see Brainerd, 1993, for reservations about the reality of stages). There have been several attempts to model this course of development using neural networks (e.g., McClelland &

Jenkins, 1991), although success has been limited hitherto (Raijmakers et al., 1996). In this paper, I will discuss an alternative approach, which uses a variant of EPAM[1] (Feigenbaum & Simon, 1984), a symbolic theory of cognition.

The chapter is organised as follows: first, I will give a brief summary of Piagetian theory. Second, I will present Feigenbaum and Simon's (1962, 1984) theory of EPAM, and will compare it to the main features of Piaget's theory; in particular I will attempt to map the concepts of assimilation and accommodation in Piaget's theory onto the concepts of familiarisation and accommodation in Feigenbaum and Simon's EPAM theory. I will then present an EPAM-like model of the balance scale task, and will discuss preliminary results on how well it accounts for children's discontinuous, stage-like development. The analysis will focus on the transition phase between rules, and will employ catastrophe theory (Thom, 1975), and in particular, the catastrophe flags proposed by Gilmore (1981; see also van der Maas & Molenaar, 1992). Finally, I will discuss the relationship between symbolic models (such as EPAM and ACT* [Anderson, 1983]) and dynamical systems.

1.1 Key mechanisms of change in Piaget's theory

Adaptation occurs whenever an organism-environment interaction modifies the organism so that its chances of survival are enhanced. At the core of Piaget's theory (e.g., Piaget, 1936) lies the assumption that two fundamental, complementary, and undissociable mechanisms—assimilation and accommodation—are necessary for an organism to adapt, both biologically and cognitively. Accommodation is the mechanism by which the organism changes its internal structure as a function of the properties of an external object. Assimilation is the mechanism by which the organism changes the object so that it fits into its own structures. It is assumed that adaptation is reached when an equilibrium is reached between the two mechanisms.

Beyond this general characterisation, the two mechanisms have never been precisely specified within Piagetian theory, which has led several authors (e.g. Flavell, 1963; Klahr, 1995) to see them as one of its main theoretical weaknesses. This is particularly damaging for the theory, since another key concept—scheme—depends directly upon these two mechanisms. Schemes, which are the structures on which assimilation and accommodation operate, "refer to *classes* of total acts, acts which are distinct from one another and yet share common features" (Flavell, 1963).

---

[1]Elementary Perceiver And Memoriser

1.2 Computer modelling of developmental transitions

Earlier information-processing accounts of Piaget's theory (Siegler, 1981; Klahr & Wallace, 1976) have used a rule-based formalism. Despite Klahr and Wallace's efforts in this direction, and despite the potential offered by cognitive architectures such as Soar (Newell, 1990) or ACT* (Anderson, 1983), the mechanisms allowing creation of new rules and transitions from old rules to new ones have never been spelled out in detail within a computational framework. The best effort in this direction is the work of McClelland and Jenkins (1991), who attempted to show how a (non-symbolic) connectionist model could simulate stage-like development using continuous learning. I will come back to this work later in the paper.

Contrary to a widely held opinion (e.g., Raijmakers et al., 1996, p. 102), however, and as will be illustrated below, there is nothing in symbolic models that make them unsuitable for learning in adaptive interaction with the environment or for implementing self-modifying systems with no fixed architecture. Actually, one of the attractive features of symbolic models (in particular production systems) is that they allow one to simulate complex dynamical systems evolving over time. In this respect, the use of symbolic models may be seen as the equivalent for cognitive science of differential and difference equations in physics (Newell & Simon, 1972). This use of computer simulations is not antithetical to the use of the mathematical theory of dynamical systems, but may be complemented by it. (See Klahr, 1995, for an insightful discussion of the differences and similarities between symbolic and non-symbolic models of cognition, in particular with respect to developmental questions).

1.3 EPAM

EPAM (Feigenbaum & Simon, 1962, 1984) is a general theory of human learning and memory, whose main goal is to formalise some of the invariants of cognition. EPAM can be visualised as "sandwiched" between a sensory-perceptual front end, which includes (parallel) mechanisms for feature extraction, and a semantic back end, which consists in semantic and procedural memory (Feigenbaum & Simon, 1984). Five postulates underlie EPAM:

1. Attentional mechanisms operate serially
2. Chunks are the basic units on which the system operates
3. Fixation of a chunk in LTM requires a constant amount of processing time per chunk (about 8 seconds)
4. Immediate memory (3 - 7 chunks) stores material temporarily

5. The central processing mechanisms fixate any part of the stimulus to which they attend. Attention is modifiable by strategies, instruction, etc.

Learning is seen as the growth of a discrimination net (DN), where nodes include the symbolic representation of objects, and where tests check the presence of features or parts of objects (see Figures 1 and 2). When perceived, an object is sorted to a node in the DN by a sequence of tests. Two cases may occur: (a) the characteristics of the external object match or partially match the internal representation (the *image*) of the node that has been reached, or (b) they mismatch it. The outcome of the match is used in two complementary learning mechanisms. *Familiarisation* adds information to a node when the internal representation matches a subset of the characteristics of the object. *Discrimination* creates a new node in the DN when there is a mismatch. Several time parameters (e.g., 8 seconds to create a new node) limit the degrees of freedom of the theory when applied to new domains.

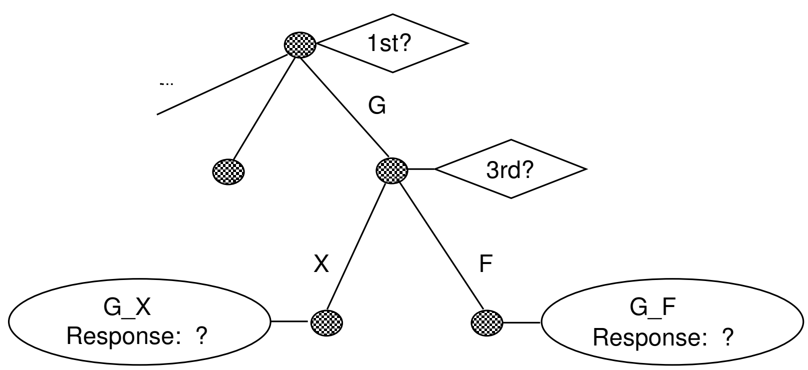

Figure 1. Example of a discrimination tree grown by EPAM in a paired-associate task, consisting in learning pairs of nonsense syllables. Shaded circles stand for nodes, diamonds for tests, and ellipses for the internal representation of a node.

One of the strengths of EPAM is that it has been able to offer detailed and often accurate simulations for a wealth of empirical phenomena in various domains: verbal learning and memory (Feigenbaum & Simon, 1984), letter perception (Richman & Simon, 1989), concept formation (Gobet, Richman, Staszewski & Simon, 1997; Richman, 1991), acquisition of syntactic categories (Gobet & Pine, 1996), and chess expertise (de Groot & Gobet, 1996; Gobet & Simon, 1996; Simon & Gilmartin, 1973).

In comparison to other cognitive architectures such as ACT*, Soar or neural nets, EPAM has the advantage of being a simple and parsimonious model, with few degrees of freedom (mainly subjects' strategies). Like other architectures, EPAM learns and makes quantitative predictions.

EPAM has features that suggest complex dynamic behaviour: it models behaviour as an adaptive interaction with the environment; its main data structure is a recursive tree, which, by the way it is constructed, has fractal properties; and, finally, its two learning mechanisms act on information in ways that are not unlike stretching and folding, which have been shown to lie at the source of chaotic behaviour in many systems (Stewart, 1989).

Recently, it has been proposed that EPAM could be extended to account for procedural and semantic memory (Gobet, 1996), by combining three learning mechanisms : (a) EPAM's perceptual learning mechanisms; (b) mechanisms connecting nodes from two nets and generating productions (condition-action pairs); and (c) mechanisms connecting nodes from three nets and generating semantic links.

Define **Epam** (Current-node, Observation)
    If Current-node is a leaf,
        Then if there is a positive difference between Observation and Current-node's image,
            Then **Discriminate** and return new image.
            Else **Familiarise** and return image.
        Else let Test be the test at the current-node,
        Find Component of Observation referred to in Test,
        If the Component is a list of components,
            Then set Component to **Epam** (Root-node, Component),
        Set Current-node to the child that corresponds to Observation's value on Test,
        Return **Epam**(Current-node, Observation).

Figure 2: The EPAM learning algorithm.

1.4 Comparison of Piaget's theory and EPAM

Admittedly, Piaget's theory and EPAM address different questions: the former is interested in development, the latter in learning. However, the similarities between these approaches are obvious, and they will be discussed here in some detail.

Accommodation implies that it is the organism that changes, not the perceived properties of the external object. Thus, accommodation covers both discrimination (creation of a new node) and familiarisation (addition of new features to the internal representation through pointers to other nodes). Accommodation also covers the creation of production and semantic links in the extended version of EPAM mentioned above. Assimilation implies that the perceived properties of the object change to fit to the structures of the organism. In EPAM, this occurs during recognition of an object, as follows: an object is recognised when a node is accessed by sorting through the discrimination net; the image (i.e. internal representation of the object) is used for further information processing, such as imagery or problem solving. Now, the properties of this image may be different from the properties of the object itself, either because the object has been misclassified or because the internal representation of the object under-represents the properties of the object. In both cases, the properties of the object have been changed as a function of the structures of the organism. Thus, these two situations correspond to Piaget's notion of assimilation. Note that EPAM provides mechanisms for what Piaget (1936) calls reproductive, recognitory, and generalising assimilation.

The presence of stages in development may be explained as follows. Conditions and actions are determined by tests, and these tests are in turn determined by biological features, for example ability to carry out a given movement, or level of development of the perceptual acuity. Changes in these biological determinants lead to a restructuring of the net, and therefore to the creation of new clusters of productions (= schemes). The same explanation can be given at a cognitive level. Changes in cognitive structures lead to changes in the way tests are carried out, which lead to the creation of new nodes through discrimination, which in turn leads to the creation of new schemes.

If a sufficiently good mapping between the two theories can be reached, this would mean that the same mechanisms lead to development and to learning. The main changes are in the way tests evolve over time. In the first case, biological maturation provokes disequilibrium, and in the second case, cognitive "maturation." In both cases, disequilibria are caused by the fact that tests that were adequate at the time of learning are not adequate anymore, because the elements constituting the tests themselves have changed, either by changes in the biological substratum (e.g. increase in visual acuity in the infant) or by changes in the cognitive substratum (i.e. increase in knowledge). Tests are just nodes in LTM, whose contents may be changed by familiarisation and whose access may be modified by discrimination.

1.5 A first look at stage-like learning in EPAM

To illustrate the dynamics of EPAM, let us consider a simple task domain, consisting in learning random strings of digits, and let us vary the probability of application of the familiarisation and discrimination mechanisms. With some of these probability values, stagewise behaviour can be obtained. Figure 3, which plots the percentage of strings correctly recognised as a function of the number of learning trials, shows a particular clear "jump" after 11 trials. However, to explore EPAM's properties in more detail, it is preferable to look at a richer domain: the balance scale task.

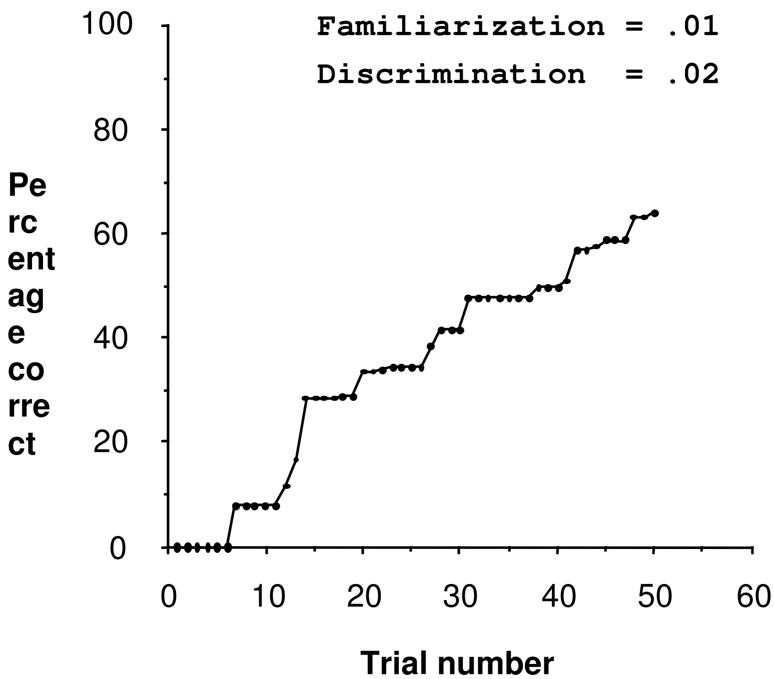

Figure 3. Stagewise learning by EPAM in a simple task.

2 Discontinuities in the balance scale task

The balance scale task, introduced by Inhelder and Piaget (1955), has been extensively studied in developmental psychology, by Siegler (1976, 1981) among others. In the standard version of the task, a balance scale is presented to the child, with some weights placed on a peg to the left of the fulcrum, and some weights placed on a peg to the right of the fulcrum. A lever prevents the balance from tipping. The child has to predict on which side, if any, the scale would go down if it were free to move.

2.1 Siegler's analysis of the task

Siegler (1976) identified a series of four rules that offer a good description of children's behaviour. Rule 1 pays attention to the number of weights on each side, but not to the distance from the fulcrum. Rule 2 is the same as Rule 1, except that it takes distance into account when the weights are the same, proposing that the side with the weights furthest from the centre will go down. Rule 3 refines Rule 2 and adds the case where weights are not equal: if the distances are equal, then the side with the most weights will go down; if the distances are not equal, then Rule 3 guesses among the three possible outcomes. Finally, Rule 4 replaces the guessing behaviour with a correct application of the torque principle (The side with the greatest Weights x Distance product will go down. If the products are equal, then the scale will stay in equilibrium).

Siegler devised six types of problem that allowed him to differentiate between the use of each rule. In *balance problems*, weights and distances are identical on both sides. In *weight problems*, distances are the same, and weights differ. In *distance problems*, weights are the same, and distances differ. In the remaining three *conflict problems*, both weights and distances differ and are in conflict (the weight is greater on one side and the distance is greater on the other). In *conflict-weight problems*, the torque is greater on the side with the greater weight; in *conflict-distance problems*, the torque is greater on the side with the greater distance; in *the conflict-balance problems*, weights and distances cancel out, and the torque is the same on both sides.

Balance and weight problems lead to 100% correct responses with all four rules. For distance problems, there is a dramatic improvement between Rule 1 (which incorrectly predicts "balance") and the other rules (always correct). Conflict-weight problems obtain 100% correct responses with all rules except Rule 3, where guessing leads to a correct answer one third of the time. Finally, with conflict-distance and conflict-balance problems, there is an improvement from Rule 1 and Rule 2 (both incorrect all the time), to Rule 3 (chance responding), to Rule 4 (always correct).

Siegler (1976, 1981) found that the behaviour of about 93% of his subjects (children older than four) could be explained in terms of one of the four rules, the criterion being that 20 out of 24 of the responses should correspond to the predictions of one of the rules. Finally, children seemed to progress through the different rules as they grew older, with the qualifications that Rule 2 is highly unstable and that Rule 4 is not always attained.

2.2 Connectionist account of stagewise development in the balance scale task

McClelland and Jenkins (1991), as well as McClelland (1995), have proposed that their connectionist model could account for the stagewise development shown by children. Recently, Raijmakers et al. (1996) have criticised this claim. Using strict criteria for the acceptance of abrupt transitions (the so-called catastrophe flags [Gilmore, 1981] derived from catastrophe theory [Thom, 1975]), they show that McClelland and Jenkins' (1991) simulations do not show abrupt transitions. In particular, the simulations of the transition between Rule 1 and Rule 2 did not satisfy three of the eight catastrophe flags: bimodality of the dependent variable, sudden jump, and non-accessibility region (the five other flags—anomalous variance, hysteresis, divergence, divergence of linear response, and critical slowing down—are typically difficult to use with psychological data; see van der Maas & Molenaar, 1992).

In the remainder of this paper, I will explore how an EPAM-like model fares in this task, focusing, as in Raijmakers et al.'s analysis of McClelland's model, on the transition between Rule 1 and Rule 2. The comparison with connectionist models is interesting, as Richman and Simon (1989) have claimed, in their study of the context effects in letter perception, that connectionist models and EPAM share important features, although they stand on different sides of the symbolic/sub-symbolic divide.

2.3 EPAM account of the balance scale task

The model consists of the juxtaposition of two discrimination nets—one net for the conditions, and one net for the actions of productions—which implements a simple production system (Gobet, 1996). The condition net encodes the perceptual representation of the external problem, while the (somewhat degenerated) action net encodes the possible responses ("left", "right", or "balance"). The discrimination and familiarisation mechanisms are as described above. As in McClelland and Jenkins (1991), it is assumed that children pay attention to the weight dimension first. In the simulations, problems where weight is the key predictor (i.e., problems with the same distance from the fulcrum on both sides) were listed ten times more often than other problems. In addition, the model pays attention to the weight attribute before paying attention to the distance attribute.

There are three parameters in the model: `fam`, the probability of carrying out familiarisation, `disc`, the probability of carrying out discrimination, and `answer`, the probability of creating a link between the condition discrimination net and one of the three nodes of the action discrimination net. These were set as follows: `fam = .15; disc = .55;` and `answer = .90`. Failure to carry out one of the three cognitive operations may reflect load on working memory, diverted attention, or time pressure from the environment. Note that with low values for these parameters, the model fails to even reach Rule 1, and with high values, the model rapidly reaches the Rule 4 level.

2.4 Training of the model

During training, the model is presented with a sequence of problems and their solution, as in the simulations with connectionist models (McClelland & Jenkins, 1991; Raijmakers et al., 1996). The training corpus was created as follows: all 625 possible problems are generated, and problems where weight is the key predictor are inserted into the corpus ten times more often than other problems.

2.5 Testing of the model

The program was tested after learning each group of 100 problems. Learning was turned off during testing. Twelve problems were randomly generated for each of Siegler's six categories (a total of 72 problems). The fit of the program to the four rules was judged using the same criterion as in Siegler's (1976) research: the program is said to "use" a rule if it gives the same response as the rule on at least 83.3% of the 72 problems.

3 Results of simulations

In general, the model is sensitive to the order of the learning problems, and different orders lead to different progressions through Siegler's four rules, some rules being skipped in some runs. In addition, there is sometimes a lot of variability from trial to trial within the same run. In general, the discrete learning mechanisms of EPAM lead to a learning curve that is less smooth than with connectionist models.

Only runs where the model approximates all the first three rules will be considered. The analysis will focus on four runs of the model. In these runs, the model fits any of the rules 78% of the time on average, which is slightly a worse fit than that reported by McClelland and Jenkins (85%).

As with McClelland and Jenkins' model, the EPAM model seems to change in a stagewise fashion when rules are plotted as a function of time. However, as correctly pointed out by Raijmakers et al. (1996), more stringent criteria are necessary to reach firm conclusions about the stagelike character of the model. Raijmakers et al.' criteria are used below to judge EPAM's behaviour. They consist in three of the catastrophe flags: bimodality of distribution, inaccessibility region, and presence of jumps.

3.1 Bimodality of distribution and inaccessibility region

Figure 4 shows the frequency distribution of the testing trials from the first time Rule 1 reaches criterion to the last time Rule 2 reaches criterion. The outcome is similar to Raijmakers et al.'s (1996) simulations of McClelland and Jenkins (1991) model, in that, contrary to what is predicted by catastrophe theory for a genuine discontinuity, there is no clear bimodality of the scores. In addition, the predicted inaccessibility region—roughly the central portion of the scores—is not present in the diagram.

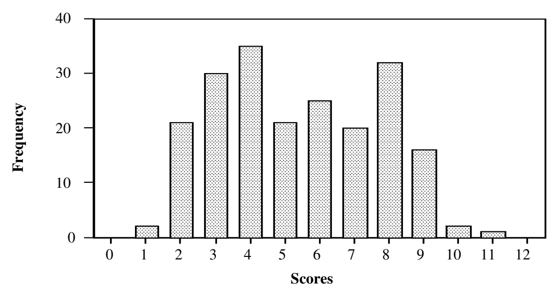

Figure 4. Frequency distribution of the scores for the testing trials from the first time Rule 1 reaches criterion to the last time Rule 2 reaches criterion.

3.2 Sudden jump

I followed a similar approach to that proposed by Raijmakers et al. (1996) to test the presence of a sudden jump. For each series of scores, the transition point is identified (first time Rule 2 reaches criterion uniquely, or third time Rule 2 reaches criterion in conjunction with Rule 1). Then, the four series are aligned on the same transition point, session 15. Presence of a jump is tested by using multiple regression with time (session number) and position with respect to possible jump (0 before, and 1 after) as independent variables. Catastrophe theory predicts that, in the case of a genuine jump, the time variable should not be significant, while the position variable should. Regression analysis shows that neither variable is a statistically significant predictor of score. The only measure that can count as evidence for a jump is the increase in variance after session 15. Note also that the scores after the transition between Rule 1 to Rule 2 are lower than those presented in McClelland and Jenkins (1991).

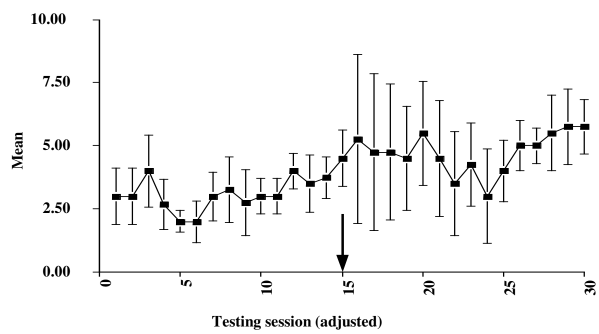

Figure 5. Mean proportion correct (four runs) in the testing phase for the trials lying between the first time Rule 1 reaches criterion to the last time Rule 2 reaches criterion. The vertical bars indicate the standard deviations. The series are shifted such that at time 15, all simulations reach Rule 2 for the first time uniquely, or for the third time in presence of Rule 1.

4 Conclusion

In this paper, I have investigated how an incremental, self-organising symbolic system could account for stagewise development. As with connectionist models, it was expected that rules would emerge from the interaction of several nodes.

The preliminary results presented above indicate that EPAM mimics rule-based behaviour, but that it does not meet the strict criteria of discontinuity derived from catastrophe theory. In this respect, EPAM behaves in a way quite similar to connectionist models. As with connectionist models, the space of possible parameters is huge, and the effects of slight changes in the learning algorithms are often unclear. Therefore, it would be premature to conclude that all EPAM or connectionist models will fail to show genuine stagewise development.

The similar behaviour of EPAM and of the neural nets studied by McClelland and Jenkins (1991) suggests that the distinction between symbolic and non-symbolic systems is not as clear-cut as is often thought. The best example of the difficulty of applying this classification is perhaps offered by Anderson's (1983) ACT* (now, ACT-R) production-system architecture, one of the most popular symbolic architectures. The ACT* architecture includes both mechanisms allowing the creation and selection of productions, and mechanisms allowing activation to be spread across nodes. Activation itself is governed by a set of differential equations (Anderson, 1983, p. 22). Thus, ACT* uses activation mechanisms that are characteristic of neural nets. But is ACT* a dynamical system? van Gelder (in press; see also in this volume) defines a dynamical system as a "set of quantitative variables changing continually, concurrently and interdependently over quantitative time in accordance with dynamical laws described by some set of equations." By this definition, the symbolic ACT* definitely qualifies as a dynamical system. (One could argue that ACT* as implemented on a digital computer lacks the continuity requirement of van Gelder's definition. But, then, so would any simulation of a dynamical system on a digital computer.) It seems important to stress that "symbolic" and "dynamical" are orthogonal properties of a system: some symbolic models can be characterised as dynamical systems, and some non-symbolic models cannot.

The symbolic system described in this chapter learns by growing a discrimination net as a function of the input from the environment. It thus embodies a self-organising, dynamical system. In addition, it shows complex behaviour characterised by non-linearities. As a consequence, it makes the view that symbolic systems are not capable of implementing self-modifying systems (e.g., Raijmakers et al., 1996) untenable. Incidentally, as shown by Vera and Simon (1993), the same conclusion applies to the view that symbolic systems cannot be situated in their environment.

At present, the idea of varying the probability of discrimination and familiarisation mechanisms is new within the EPAM framework, and the underlying dynamics are still poorly understood. Future work should establish whether true stage-like discontinuity can emerge from such an approach, or whether additional mechanisms need to be added, such as the creation of semantic links between the nodes (Gobet, 1996) or the creation of templates (Gobet & Simon, 1996), which are data structures allowing the storage of variable values.

Reference list

1. Anderson, J. R. (1983). *The architecture of cognition*. Harvard University Press, Cambridge, MA.
2. Brainerd, C. J. (1993). Cognitive development is abrupt (but not stage-like). Commentary of Thomas, H., & Lohaus, A. (1993). Modeling growth and individual differences in spatial tasks. *Monographs of the Society for Research in Child Development, 58*, (9, Serial No 237).
3. de Groot, A. & Gobet, F. (1996). *Perception and memory in chess. Heuristics of the professional eye.* Assen: Van Gorcum.
4. Feigenbaum, E . A. & Simon, H. A. (1962). A theory of the serial position effect. *British Journal of Psychology, 53*, 307-320.
5. Feigenbaum, E. A., & Simon, H. A. (1984). EPAM-like models of recognition and learning. *Cognitive Science*, *8*, 305-336.
6. Flavell, J. H. (1963). *The developmental psychology of Jean Piaget.* Princeton, NJ: Van Nostrand.
7. Gilmore, R. (1981). *Catastrophe theory for scientists and engineers.* New York, Wiley.
8. Gobet, F. (1996). *Discrimination nets, production systems and semantic networks: Elements of a unified framework.* Proceedings of the Second International Conference of the Learning Sciences.
9. Gobet, F., & Pine, J. (1997). Modelling the acquisition of syntactic categories. *Proceedings of the 19th Annual Meeting of the Cognitive Science Society*, p. 265-270. Hillsdale, NJ: Erlbaum.
10. Gobet, F., Richman, H., Staszewski, J., & Simon, H. A. (1997). Goals, representations, and strategies in a concept attainment task: The EPAM model. *The Psychology of Learning and Motivation, 37,* 265-290.
11. Gobet, F. & Simon, H. A. (1996). Templates in chess memory: A mechanism for recalling several boards. *Cognitive Psychology, 31,* 1-40.
12. Inhelder, B. & Piaget, J. (1955). *De la logique de l'enfant à la logique de l'adolescent.* Paris: Presses Universitaires de France.
13. Klahr, D. (1995). Computational models of cognitive change: The state of the art. In T. J. Simon, & G. S. Halford (Eds.), *Developing cognitive competence: New approaches to process modeling*. Hillsdale, NJ: Erlbaum.
14. Klahr, D. & Wallace, J. G. (1976). *Cognitive development: An information processing view.* Hillsdale, NJ: Erlbaum.

15. McClelland, J. L. (1995). A connectionist perspective on knowledge and development. In T. J. Simon, & G. S. Halford (Eds.), *Developing cognitive competence: New approaches to process modeling*. Hillsdale, NJ: Erlbaum.
16. McClelland, J. L., & Jenkins, E. (1991). Nature, nurture and connections: Implications of connectionist models for cognitive development. In K. VanLehn, (Ed.), *Architectures for intelligence*. Hillsdale, NJ: Erlbaum.
17. Newell, A. (1990). *Unified theories of cognition*. Cambridge, MA: Harvard University Press.
18. Newell, A., & Simon, H. A. (1972). *Human problem solving*. Englewood Cliffs, NJ: Prentice-Hall.
19. Piaget, J. (1936). *La naissance de l'intelligence.* Neuchâtel: Delachaux et Niestlé.
20. Raijmakers, M., van Koten, S., & Molenaar, P. (1996). On the validity of simulating stagewise development by means of PDP networks: Application of catastrophe analysis and an experimental test of rule-like network performance. *Cognitive Science, 20,* 101-136.
21. Richman, H. B. (1991). discrimination net models of concept formation. In D. H. Fisher, M. J. Pazzani & P. Langley: *Concept formation: Knowledge and experience in unsupervised learning*. San Mateo: Morgan Kaufmann.
22. Richman, H. B., & Simon, H. A. (1989). Context effects in letter perception: Comparison of two theories. *Psychological Review*, *3*, 417-432.
23. Siegler, R. (1976). Three aspects of cognitive development. *Cognitive Psychology, 4*, 481-520.
24. Siegler, R. S. (1981). Developmental sequences within and between concepts. *Monographs of the Society for Research in Child Development, 46* (2, Serial No 189).
25. Simon, H. A., & Gilmartin, K. J. (1973). A simulation of memory for chess positions. *Cognitive Psychology*, 5, 29-46.
26. Stewart, I. (1989). *Does God play dice?* Cambridge: Blackwell.
27. Thom, R. (1975). *Structural stability and morphogenesis.* Reading, MA: Benjamin.
28. Thomas, H., & Lohaus, A. (1993). Modeling growth and individual differences in spatial tasks. *Monographs of the Society for Research in Child Development, 58*, (9, Serial, No 237).
29. van der Maas & Molenaar (1992). Stagewise cognitive development: An application of catastrophe theory. *Psychological Review, 3*, 395-417.
30. van Gelder, T. (in press). Dynamical approaches to cognition. *MIT Encyclopedia of the Cognitive Sciences.* Cambridge, MA: The MIT Press.
31. Vera, A.H., & Simon, H.A. (1993). Situated action: A symbolic interpretation. *Cognitive Science, 17*, 7-48.