

Production Systems for Children's
Seriation Behaviour: an Introduction
and an Example

Richard M. Young

Carnegie-Mellon University

Department of Psychology

Report #71-31

Abstract

A production system (PS) is a set of rules each of the form $C \Rightarrow A$, meaning that the action A is to be taken whenever the condition C is satisfied. PSs can be used to construct information-processing models of human behaviour, and have certain advantages over other simulation techniques that make them especially suitable for describing cognitive development. In this paper the nature of children's seriation behaviour is discussed, and an illustrative PS is written to model the actual behaviour of a 5-year-old subject. A technique for the experimental validation of PSs is outlined.

Preface and Acknowledgements

This paper is essentially a re-written version of my dissertation proposal, and aspires to be a self-contained introduction to the use of productions systems for modelling problem solving in general and children's seriation behaviour in particular. The main changes I have made are to omit the less relevant parts of the literature survey and (some of) the wilder speculation, and to add a couple of pages discussing the relationship between production system analysis and the techniques of the developmental psycholinguists.

I have some reservations about the example presented in Section 3. S315's protocol was originally analysed over eighteen months ago, and since then my ideas have changed enough that the analysis would look rather different were I to re-do it now. I have decided to include it anyway mainly because its (unrealistic) simplicity makes it an especially easy-to-follow illustration of the production system approach. But the reader should not take it as reflecting my current notions either about the nature of seriation behaviour nor of the way a production system should be written.

I would like to thank the members of my dissertation committee, especially Dr. Allen Newell, and the several other persons who made useful comments on an earlier version of this paper. I would also like to thank Dr. Lauren Resnick for making available the videotape from which the S315 protocol is drawn.

This research was supported by Public Health Service Research Grant MH-07722 from the National Institute of Mental Health.

TABLE OF CONTENTS

	Page
Abstract	i
Preface and Acknowledgements	ii iii
Table of Contents	iv
1. Production Systems and Seriation Behaviour	1
2. Seriation	17
3. A Detailed Example: S315 on the Seriation Task	29
4. Validation of Production Systems	39
References	45

1. Production Systems and Seriation Behaviour

The relationship between computers and psychological theorising has been changing in recent years. At first, computer simulations served mainly to counter certain objections to the language of cognitive theories. To many psychologists it seemed that explanations phrased in terms of internal representations and cognitive maps presupposed the existence of someone to examine the representation or read the map, and to postulate a mysterious homunculus in order to account for one's behaviour was clearly to beg the issue. The advent of computers soon laid this particular ghost to rest. Computer scientists began to speak of their charges as following rules, making decisions, and so on, but since "the computer was clearly a mechanistic device, behaving as it did because the laws of physics constrained it to" (Newell & Simon, 1971), there was no temptation to level the same accusation of anthropomorphism. It became accepted that if dumb mechanism can perform these activities without the assistance of something or someone else inside actually to do the work, then so can the human brain. This exorcism of the homunculus-as-executor perhaps accounts for the popularity during the 1960's of computer programs to simulate psychological processes, from LT and TOTE units onwards (Newell, Shaw & Simon, 1958; Miller, Galanter & Pribram, 1960).

For a demonstration of sufficiency of this kind, the internal details of the simulation program are relatively unimportant. It matters little what language the program is written in, nor how it is structured. What counts is that the program indeed perform the task it claims to be simulating. Kessen (1962) provides an extreme statement of this view when he says that

"Given a set of circumstances and a set of responses, the problem of the simulator is to write a machine program which will give him the answers the human being gave. In devising a program he may be neutral with respect to general psychological theory, however unlikely this case is in fact, and he may, if he chooses, ignore the problem of transition from one stage to another or the problem of differences among species or other sub-populations. His task is to make guesses about the functioning of the 'black box' and then, in imitation of Providence, to build it."

Recently, however, this view has been challenged; some have objected to the idea that arbitrary computer programs are suitable as theories of human behaviour. Fodor (1968), for instance, has argued that in order to constitute a valid simulation, it is not enough for the program simply to display the same behaviour as the human did. It must also be capable of, but refrain from, displaying the whole range of potential behaviour that the human could have displayed, but didn't. (Fodor himself expresses the point a little less laconically.) Then there is the matter of the supposed origin of the program. It happens that computer programming is a specialised, skilled intellectual activity, and for many people a hard one to learn. Indeed, it is sufficiently difficult and poorly understood that attempts to mechanise it remain at a rudimentary level, in spite of the effort being applied to the task (for example, Waldinger & Lee, 1969). How, then, can it make sense to speak of a person's following a program inside his head: how could the program have got there? Where did it come from? It seems implausible that the person himself could have devised it, since in most cases the programming task is well beyond his (or anyone else's) ability. It appears that in using programs to explain behaviour we are again in danger of invoking a homunculus, this time in the role of programmer. Or again, Frijda (1967) has pointed out that it is not sufficient to judge a program that

claims to be a simulation of human behaviour simply by examining its external performance. One must also examine its internal structure. As Earl Hunt (1971) puts it,

"... very large programs have a way of becoming bogged down in details. It may be that it will be hard to find out what causes the simulation to behave in a certain way. Of course, we can claim that if we have a running program, we must have a formal model, but this is not enough. The model is supposed to aid our understanding, not increase our confusion. We already have one black box, man, and hardly need another made of IBM cards."

The last few years have seen progress in the development of a new technique which holds promise for answering both the complaint of arbitrariness and the question of origin, thereby banishing also the homunculus-as-programmer. Devised in order to help in the analysis of detailed protocols of human problem solving behaviour, the technique attempts to capture the regularities in a subject's behaviour by writing a set of independent rules -- known as a production system (PS) -- that expresses what the subject does under what conditions (Newell, 1968; Newell & Simon, 1971). Each individual rule (or PR, for "production rule") is a condition-action statement of the form $C \Rightarrow A$, and means simply that in the circumstances specified by C the subject performs action(s) A. As a simple example of a PS, Newell & Simon (1971) give the following PRs to describe the behaviour of a thermostatic control system intended to keep the temperature of a room between 70° and 72° :

- (T1) Temperature $< 70^{\circ}$ and Furnace-state = off $\Rightarrow$ Turn-on [furnace]
- (T2) Temperature $> 72^{\circ}$ and Furnace-state = on $\Rightarrow$ Turn-off [furnace].

The action on the right hand side of T1 applies whenever the condition on its left hand side is satisfied, and similarly for T2. An extended discussion of the properties of PSs, and more elaborate examples, can be found in the references given above.

To date, PSs have been applied almost exclusively to the analysis of adult verbal protocols drawn from domains such as chess, formal logic, and cryptarithmic. The research this paper is based on attempts to break new ground^{1,2} in a number of ways:

- it uses PSs to account for the non-verbal behaviour of children;
- it is less closely tied to protocol analysis as its principle means of empirical validation;
- and it confronts seriously the question of how PSs change over time.

In this paper we will argue the case for PSs and show how they are useful in trying to understand children's behaviour on a seriation task. We will discuss certain properties of PSs and their advantages over other forms of simulation (Section 1), review briefly what is known about the nature of seriation behaviour (Section 2), and then write a PS to model the actual problem-solving behaviour of a young subject (Section 3). Finally, in Section 4, an experimental approach to the validation of PSs is outlined. However, we begin by making the acquaintance of a hypothetical PS in order to have some idea of the kind of beast we are dealing with.

¹ Gascon (1969) has examined behavioural protocols of children performing a weight seriation task, but his "production system" forms only a small part of his simulation program.

² I have only recently become aware of the interesting work of Knifong (1971), who carried out a "framework analysis" of detailed protocols of young children working with a beam balance. Without pretending to do justice to the merits of his technique, we can perhaps characterise it as complementary to the approach adopted here. A "framework" is essentially a specific belief the child has about the topic being investigated (i.e. in Knifong's case, the operation of the beam balance; in ours, the seriation task). It accounts for the details of the child's behaviour in the sense that the behaviour has to be consistent with it and fall within its scope. Production systems, on the other hand, try to model the details of the behaviour more directly, but perhaps stand in need of something like a "framework" to represent explicitly the knowledge implicit in the production rules.

The seriation task and an idealised production system

The particular problem-solving task we will be concerned with in this paper is quite straightforward: a child is given half-a-dozen wooden rods of different lengths and asked to arrange them in a straight line in order of size. We consider only cases where the child moves the rods one at a time between the initial heap -- called the Pool -- and the geographically distinct series he is constructing -- called the Line. The hand holding the rod will be referred to as the Hand, though it may sometimes be the child's left hand and sometimes his right. We concentrate on the repetitive, cyclic phase of the seriation, which usually involves selecting rods from the Pool one by one and placing them in the Line. For present purposes we will not even attempt to model accurately the start and termination of the seriation process.

Consider a hypothetical child who solves the problem by always picking up the smallest rod remaining in the Pool, and then placing it at one end of the growing Line. Ignoring for the moment all questions of verification or psychological plausibility, the PS below can be used to simulate his behaviour. The left hand side of each of the PRs is referred to as its condition, and consists of a goal (Place.all, Place, or Get.next), possibly conjoined with a predicate (At, or "Hand ="). The right hand side, or action, consists of an operator (To, Set.goal, Select.smallest, or Insert.at.right). The rules apply one at a time as determined by the immediate goal and the states of the predicates: at any moment the PR whose condition is satisfied is the next to be invoked. Usually the rules fire in the cyclic sequence P0; P1, P5, P3, P2, P6, P4; P1, P5, P3, P2, P6, P4; Rules P1, P3, and P5 can be seen to carry out the task of getting the rods, and rules P2, P4, and P6 the task of placing them. The 'translation' given with each PR should make the PS almost self-explanatory.

(P1) Place.all; Hand = 0 $\Rightarrow$ Set.goal [Get.next]

If the immediate goal is to place all the rods (i.e. to solve the seriation problem), and if the Hand is empty, then set up the goal of getting the next rod.

(P2) Place.all; Hand = 1 $\Rightarrow$ Set.goal [Place]

If the immediate goal is to place all the rods, and the Hand is holding a rod, then set up the goal of placing it correctly.

(P3) Get.next; At [Pool] $\Rightarrow$ Select.smallest

If the immediate goal is to get the next rod, and the Hand is at the Pool, then pick up the smallest one.

(P4) Place; At [Line] $\Rightarrow$ Insert.at.right

If the immediate goal is to place the rod, and the Hand is at the Line, then put the rod at the right hand end of the growing series.

(P5) Get.next $\Rightarrow$ To [Pool]

(P6) Place $\Rightarrow$ To [Line]

If the immediate goal is to get the next rod, (and the Hand is not at the Pool), then move it to the Pool. Similarly, if the goal is to place the rod (and the Hand is not at the Line), move it to the Line. P5 and P6 do not apply if P3 and P4 are applicable, respectively -- see the discussion below.

Finally we need a rule to put the child in context and set him working on this task rather than some other one. Without attempting to spell out the details, we can write this as

(P0) Seriation.instructions $\Rightarrow$ Set.goal [Place.all]

i.e. when given appropriate instructions and materials, set up the goal of solving the seriation problem.

A few fairly technical remarks on this PS are needed. First, the order of the PRs. Unlike the PSs used by Waterman (1970) and Newell & Simon (1971) our PRs are effectively unordered. That is, the order in which the rules are written conveys no information about their operation. Instead of using ordering to resolve the conflict between two or more PRs whose conditions become satisfied simultaneously, we will always pick the PR whose condition is most restrictive. Thus P3, whose condition is "Get.next; At [Pool]", may conflict with P5, whose condition is "Get.next", in which case P3 will be evoked. Similarly, P4 takes precedence over P6.

Second, goals and the goal mechanism. This paper is not the place to discuss in detail what is meant and implied by the use of goals; that has been done elsewhere -- for example, in Chapter 14 of Newell & Simon (1971). For present purposes a goal can be regarded as an internal symbol whose function is to trigger behaviour that tends to bring about some state of affairs. Behaviour -- that is, a PR -- is evoked always in the context of some goal, called the immediate goal (IG). Whenever a goal is set up, the old IG is "pushed down", and the new goal is stacked above it to become the new IG. When that new goal is satisfied it is "popped off" the stack and the previous IG is reinstated. In the PS given above the stack of goals never becomes deeper than two. The goal of Get.next is satisfied by the performance of Select.smallest, and Place is satisfied by Insert.at.right. The main goal of Place.all, presumably, is satisfied by the successful completion of the seriation task.

Lastly, the matter of what Newell & Simon (1971) refer to as the parameterisation of the production rules. It often happens that one or more PRs can be regarded as versions of a more general rule specialised to the particular case of seriation (or whatever the task happens to be). Both P5 and

P6 above are essentially instances of a rule something like "if you want to do anything that requires being at a certain place, then go there". The issue is then at what level it is appropriate to represent a PR: should it be represented as a parameterised instance of the most general rule possible, or as if it existed independently as a fully specified variant? For the moment we will simply write the rules in whatever form seems most suitable. But it may be helpful to remember that the same PR can usually occur in different versions at different levels of generality.

Production systems for cognitive development

Production systems have been used in the past almost exclusively to model problem-solvers assumed to remain constant over a period of time; almost none has dealt with the possibility that the PS might change during the course of the task. The one exception we know of is Waterman's (1970) PS to play draw poker that learns to improve its game by modifying its existing PRs and adding new ones. In this section we will argue that PSs provide an excellent tool for modelling also the growth and development of cognitive skills. The arguments for their superiority over ordinary programs will be confined to issues directly related to psychology--there exist further arguments for their superiority on technical grounds, but they belong in another paper. The discussion that follows draws heavily on the eight points mentioned in Chapter 14 of Newell & Simon (1971).

An attractive feature of a PS is that its simple organisation as a homogeneous set of rules leads to a corresponding structure in the resulting behaviour. Each PR represents a fragment of potential activity--a meaningful component of the total problem-solving process. In the simple thermostat

above, rule T1 has a unique interpretation quite independent of what other rule(s) may be present; namely, "if the temperature is less than what it should be, make sure the furnace is turned on." Similarly, mutatis mutandis, for T2. This property makes it possible for a PS to be expanded incrementally, simply by adding further PRs to the existing ones. For example, Newell & Simon's simple thermostat has a more advanced feature to take special action when the temperature falls too low. We simply adjoin to T1 and T2 the new rule:

(T3) Temperature $\leq 32^{\circ}$ $\Rightarrow$ Call-repair-man; Turn-on [Electric-heater] or whatever.


Even though production systems are a form of programming language and a PS constitutes a bona fide information-processing model of the behaviour it simulates, there are important distinctions between a PS and a program written in one of the more familiar languages such as Algol, Fortran, or Snobol. (We will refer to such programs as "flowchart-programs".) These distinctions indicate that the approach described in this paper differs markedly from earlier attempts to devise information-processing accounts of cognitive development (e.g. Gyr & Fleischer, 1967; Klahr & Wallace, 1970). Klahr and Wallace, for instance, view the child's performance on a set of classification tasks as being governed by the assembly of certain subroutines into a full program. But this form of self-programming, and the manipulation of subroutines inside their heads, would seem to be just the kind of thing that children are unable to do, and more recently Klahr and Wallace (1972) have shifted to using a PS to describe their behaviour (in this case, on a class inclusion problem). Subroutines are just not manipulable in the way that PSs are. In particular, it is generally difficult to 'grow' a working

subroutine by making small, localised additions. Since it is designed to perform a logically unitary and well-defined job, one tends to think of a subroutine as having an all-or-none existence: it is either present, in which case all of it must be present, or else absent, in which case there is no point in having various pieces of it lying around. Subroutines do not share the incremental nature of production systems.

A further advantage of PSs over flowchart-programs is that it is more meaningful to ask the question "Where does it come from?" of a single PR than of a component of a typical subroutine. A subroutine usually makes sense only when viewed as an integrated whole, so the typical justification given for one of its components is the "conservative" one that when the subroutine is put together the way it is, it does the job it is supposed to. In contrast, individual PRs frequently possess a kind of 'local plausibility' that serves to justify them in their own right, regardless of the task they happen to perform when put together. This makes it easier to believe that the child could have acquired them without benefit of any deep insight into the particular task he uses them for. We pointed out above that rules P5 and P6 of the PS for seriation are versions of a more general rule expressing the idea that if you need to be somewhere, the appropriate behaviour is to move there; and it is quite easy to accept that rule as being just the kind of thing the child is likely to have learned from the explorations of his sensori-motor period of development (Piaget, 1954).

According to Newell & Simon (1971), "production systems are the most homogeneous form of programming organisation known". One result of this homogeneity is that a PS, unlike a flowchart-program, makes clear exactly which decisions the child has control over and which others are forced on

him by the environment in which he operates. To see this, consider first a fragment of a flowchart program of the form



that simulates a child performing action A followed by action B. Why are the actions done in the order A then B instead of the other way round? The flowchart provides no indication whether the sequence is decided by the child's problem-solving processes, in the sense that B could equally well be carried out before A as after it, or whether the order is determined by the very nature of the task, as would be the case if A were the action "pick up the doll" and B were "squeeze the doll". Similarly, granted that A occurs first, why is it followed by B rather than, say, C: is it because the fact of having just completed A itself serves as a cue to do B, so that the sequencing is accomplished essentially by what Osgood (1953) calls "response-produced cues"? Or because one effect of performing A is to create conditions that make B appropriate, in which case B would be next even if those conditions were set up by something other than A? The flowchart fails to say.

Production systems, on the other hand, make explicit distinctions that provide answers to these questions. In our idealised PS for seriation, we mentioned that P6 is invoked immediately before P4. However the PS contains no representation of a decision, say, not to perform P6 after P4. Rather, the order is determined by the physical impossibility of placing a rod in the Line without first moving the Hand there. Similarly, P3 is invoked when it is, not because it follows P5 but because picking up the smallest rod is the appropriate thing to do when one's Hand is at the Pool and

one wants the next rod. P3 would be invoked in these conditions no matter how they came about. Indeed, initially the child's Hand may very well already be at the Pool, in which case P5 will not be invoked the first time around and P3 will directly follow P1. It is worth looking a little more closely at this dependence upon the physical situation.

The role of the environment

One of the dimensions along which human problem solving can vary is the extent to which it is stimulus-independent rather than stimulus-bound. Examples of both kinds are easy to find. When someone is given a mental arithmetic task, say to multiply a three-digit by a two-digit number in his head, once he has the two numbers all further data are generated internally by his own mental processing and his activity becomes largely stimulus-independent. Conversely, in a perceptual-motor task such as rolling a steel ball through a maze, the subject is continually taking in information from the changing physical situation, and that new information largely determines his actions. We say that his problem-solving is highly stimulus-bound. The point we were making above can be restated in terms of this distinction. In a flowchart-program the normal flow-of-control is assumed to be data-independent, making it unsuitable for simulating behaviour that ranges widely between stimulus-independent and -bound. In contrast, a PS can deal quite uniformly with any point on the dimension, since the condition on the left hand side of a PR can refer either to an internally generated symbol (such as a goal), yielding stimulus-independent activity; or to the external situation, yielding stimulus-bound activity; or to a mixture of the two.

Thus a PS can represent the role of the environment in governing the child's behaviour in a way that a flowchart program cannot. This is important

because a pre-operational child relies to a great extent on the physical environment as a place-keeper and external memory to reduce his cognitive load (Bruner, Olver & Greenfield, 1966). Accordingly, it seems unprofitable to attempt to model cognitive development with programs that are largely stimulus-independent.

A closely related point is that it is the PS together with the properties of the environment that determine the subject's behaviour, not (except in extreme cases) the PS alone. Consider again the example of the simple thermostat. Although rules T1 and T2 describe the properties of the control system, they are not sufficient for us to predict its behaviour. For instance, at some moment T1 is invoked and the furnace is turned on, but the PS omits to tell us how long it will be before the temperature reaches 72° and T2 applies. Suppose, however, we told that when the furnace is on the temperature rises by 1° per minute and when it is off it falls by $\frac{1}{2}^{\circ}$ per minute. Then we can predict the transitions of the control system, minute by minute, starting from any initial temperature.

One can see that a PS, like any partially data-dependent process, is essentially incomplete. Since the behaviour of the PS depends upon properties of the task situation, in order to predict the actual course of a problem-solving process we must add to the PS an equally explicit characterisation of the external physical environment. From the point of view of the PS, this characterisation has to describe what effect each of the possible operators has on the states of the predicates. For example, in the case of the idealised production system for seriation given earlier, the characterisation of the environment has to include the fact that picking up a rod results in there being a rod in the Hand, i.e. that performing the operator `Select.smallest`

results in the predicate "Hand = 1" being true. Knowing this, we can predict that P2 will be the next rule to be invoked after P3. Similarly, there has to be a statement of the fact that the operator Insert.at.right leaves the Hand empty; and also four sub-statements to express that if one moves to the Pool, the Hand is then "At" the Pool, and is not "At" the Line; and so on.

Notice that this characterisation provides a partial description of the environment in terms of exactly the same predicates and operators as appear in the PS. One can readily see that a more complicated PS with more predicates and operators would require a correspondingly more elaborate characterisation to describe its environment. This point is crucial to understanding where new PRs come from, which is an issue we will not pursue here. The same point also explains one sense in which we can speak of there being a "match" between a child's abilities at any time and the situation in which he operates (see J. McV. Hunt, 1961, 1969). The fact that a different PS requires a different complementary description of the environment demonstrates that the characterisation is not at all a neutral summary of the external environment. Rather, it describes the effective environment-for-the-child in which his problem-solving takes place. This notion that different children (or the same child at different times) operate in different worlds should sound familiar. It is closely related to the observation of Newell & Simon (1971) that someone working on cryptarithmic, for example, who can make inferences such as "odd + odd = even" is in a quite different problem-space to someone who cannot.

Production systems and developmental psycholinguistics

Our use of production systems to describe the seriation abilities of children of various ages is reminiscent of the work of the developmental

psycholinguists who seek to capture a child's linguistic growth in a series of grammars (see, for example: Bellugi & Brown, 1964; Chomsky, 1969; Klima & Bellugi, 1966; McNeill, 1966, 1970). Just how similar are these two lines of research? The brief discussion below fails to do justice to the subtlety of some of their relationships, but would seem to indicate that they are sufficiently different for no useful purpose to be served by blurring the distinctions between them.

The similarities are certainly there. We find the counterpart of Klahr and Wallace's (1970) notion of 'developmental tractability' in the attempts of the psycholinguists to model the child's linguistic development by evolving systems of rules, each system emerging as a modification of the preceding one. For example, McNeill (1970) sees the child forming classes of words by successive differentiations of earlier classes, and describes the grammar as evolving in response to the pressure from these newly formed categories. Menyuk (1969) also stresses the continuity between successive stages, discussing various versions of the base-structure rules in terms of their growth from the previous stage; and Bloom (1970) similarly makes detailed comparisons between the rules of consecutive stages. Furthermore, the similarities are not restricted to matters of formalism. Both protocol analysis and developmental psycholinguistics are marked by the attention they pay to the details of an individual subject's behaviour, and their attempt to account for it. And Bloom's (1970) emphasis on the importance of understanding what the child is saying as a step in revealing his linguistic competence is made in the spirit of a born protocol analyser.

Despite these correspondences, to regard the two lines of research as essentially similar is to be misled by the trivial coincidence that both

individual PRs and grammatical rules are written with a right arrow. A PS and a grammar, after all, are quite different kinds of beast: one describes a process, and the other a pattern. With a PS, one is interested in the step by step activity of the system as the rules are invoked, each leading to some (internal or external) behaviour. With a grammar, on the other hand, one is interested only in the final outcome: the (possibly infinite) complete set of sentences that the grammar encompasses. A child's grammar does not aim to predict what he will actually say on any given occasion. Rather, it summarises his potential, and provides a description of the whole class of sentences he could utter given appropriate circumstances. This is just the difference between 'linguistic competence' and 'verbal behaviour'.

The fact that the similarities of their techniques are misleading still leaves open the possibility of a much deeper connection between the two disciplines: namely, that a person's linguistic abilities may be mediated by a mechanism that can best be described by a production system. But in our present stage of ignorance that line of thought yields only speculation, and leads far outside our proper concerns. It is high time to turn to the subject of seriation.

2. Seriation

Seriation was introduced into the psychological literature by Piaget, two of whose books report studies dealing with the task (Piaget, 1952; Inhelder & Piaget, 1964). For Piaget, the ability to seriate is of great importance for the child and underlies many of the other operations of concrete thinking; for instance, seriation and classification together provide the twin supports on which the child's (or the adult's) conception of number is based. According to Furth, "Piaget's aim in observing seriation behavior is first and foremost to discover the general behavior of relating and ordering that can become manifest in a great variety of different concrete situations" (Furth, 1969, p. 15).

Piaget's original observations on Genevan children have been replicated on children from other cultures, such as by Elkind (1964) in the U.S.A. and Proscura (1969) in the Ukraine, and even on retarded children (McManis, 1968, 1970). It is generally agreed that the ability to seriate emerges later than the conservation of length, but before the ability to make inferences based on transitivity (Elkind, 1964; Murray & Youniss, 1968; McManis, 1970; but see also Battisti & Simmons, 1968). But beyond the replications and clarifications, little work has been done. Particularly noticeable is the lack of any systematic study of the developmental relations between seriation and a range of other abilities, along the lines of Smedslund's (1964) investigation of transitivity and conservation or the scalogram analyses provided by Wohlwill (1960) for the concept of number and by Kofsky (1966) for classification. The experimental studies consist mainly of parametric variations of Piaget's original observations, as represented by Shantz's (1967) demonstration

that adding irrelevant information makes the seriation task more difficult, the effect diminishing with age. These empirical findings seem to have added little to our understanding of seriation, and none of them has been able to isolate the processes underlying successful performance of the task.

The nature of seriation

Although nothing emerges from the published studies that can reasonably be called a theory of seriation, we will see that we can usefully sharpen our intuitions about the task. In this section we will try to understand what it is that the pre-operational child is doing when he is seriating. One point crucial to appreciating the kind of problem that seriation poses for the child is to bear in mind that seriation is a task in concrete reasoning usually given to children just about to enter their stage of concrete operations. Flavell (1963, Chapter 4) characterises this stage as one in which the child's cognitive activity takes the form of mental experiments carried out on physically present objects. That is, he can "run off reality sequences in his head" without actually manipulating the objects, but he is unable yet to reason about them abstractly. Glick & Wapner (1968) have demonstrated experimentally the great difference that exists for the child between dealing with manipulable objects and simply reasoning about them. They faced the child with a problem in transitive inference; in our case, we have to remember that the child is being asked actually to seriate some real blocks. His behaviour would be quite different on a 'formal seriation' task, such as being asked: "If I were to give you a jumbled heap of sticks all of different lengths, how would you set about arranging them in a row in order of increasing size?"

Up to a certain age (Piaget's stage III) even children who can seriate correctly do so by trial and error. Although seriation of this sort appears early if the materials are easy enough--children of 18 months can typically arrange a set of nesting boxes or build a tower with blocks of decreasing size--"as soon as the size differences are made less apparent, the child manifests an almost total lack of seriation ability"(Furth, 1969, p. 64). Rather than building the line systematically, say by working step by step from smallest to largest, young children simply try rearranging the blocks in a more or less effective way until they look right. Often they make some attempt to work from smallest to largest, but they produce reversals and omissions, and the blocks left over are simply placed at either end. Failures to seriate commonly show up as a collection of subseries, instead of one integrated seriation.

This behaviour can be interpreted as showing that the child is doing little 'reasoning' about the task, and indeed he shows little anticipation of what the results of his actions will be (Inhelder & Piaget, 1964). Elkind (1964) provides the most insightful account of this stage of empirical seriation. As he interprets Piaget, the child at this stage does not represent the seriated line as part of a series that in principle can be extended indefinitely, as does the adult. Instead he views it as a particular configuration to be reproduced. Elkind compares this view of seriation to the task of fitting together a jigsaw puzzle. Structuring the seriation task as one of reconstructing an arbitrary, given arrangement of pieces inevitably leads the child to adopt a strategy of trial-and-error arrangement, and contrasts with the adult's structuring of the problem as one of exhibiting a rank ordering.

Since at this stage of trial-and-error seriation the child is concentrating on the configuration as a whole, the actual length of the rods is relevant only in so far as it affects the appearance of a rod relative to the others. As we have already seen, a child who cannot seriate operationally also does not yet possess the notion of the transitivity of length--the operational knowledge that if rod D is longer than rod C ($D > C$), and $C > B$, then necessarily $D > B$ without there being any need to compare the two directly. Inhelder & Piaget (1964) use the term "global" seriation to refer to the child's preoccupation with the overall shape of the configuration at the expense of paying attention to the elements of which it is composed. They contrast it to the operational "analytic" approach, and demonstrate that the child can anticipate the global form of the series--by drawing it as a sequence of rods of monotonically decreasing length--long before he can predict where each rod will appear in the final series.

At stage III, once he has attained reversibility, the child reaches an equilibrium in which he is able to structure the series in an operational way and follow both its directions at once. Each rod is understood to be at the same time larger than one of its neighbours and smaller than the other. Thus an understanding of the relations ' $<$ ' and ' $>$ ' lies at the heart of operational seriation. Prior to this the observation of $C < D$ led to the characterisation of C as 'small' and D as 'big', so that the simultaneous perception of the relations $B < C$ and $C < D$, far from being combined into the compound relation $B < C < D$, actually led to a contradiction: 'C is big' and 'C is small'. However, once the relations are mastered, the whole series $A < B < C < D < E \dots$ can be seen as composed of a repetition of the single relation ' $<$ '. Or, to express the child's understanding more precisely, since he functions at

the concrete level rather than the formal, we should say that he represents


the series  as a repetition of the relation . In this way, the series is comprehended operationally by the coordination of

its extensive () and intensive () aspects. This same coordination of intension and extension is what leads the child to a successful grasp of hierarchic classification, which follows close on the heels of seriation (Inhelder & Piaget, 1964).

Piaget takes as a criterion of operability the child's ability to insert an extra rod into an already formed series (Piaget, 1952). At the pre-operational stages the child is usually unable to do this correctly. Typically he places the new rod at one end of the line, or next to it, or even on top of it; or he may be able to replace one of the original rods with the new one, and then have a different one left over. Again, Elkind's analogy to a jigsaw puzzle is helpful, and suggests that the child may not so much misunderstand the task as regard it as arbitrary. Imagine completing a jigsaw puzzle only to have someone hand you another piece with instructions to "place this where it belongs!" It may also be that because the rods are seen as belonging to a closed configuration, it is hard for the child to compare the new rod against the individual elements in the series. At any rate, mastering the relations ' $<$ ' and ' $>$ ' makes the task of insertion identical to the original task of seriation. To insert a new rod D' between D and E is to manifest simultaneously the relations $D < D'$ and $D' < E$, just as D itself was placed so as to realise $C < D$ and $D < E$.

Perception, language, and seriation

It is rather tempting to attribute the attainment of seriation to the child's perceptual development; in other words, to try to explain seriation in terms of perception. Certainly in the preceding discussion we have been emphasising that the child at stage I does not structure the line as an ordered series, whereas the operational child does. But Piaget and his followers argue that to treat the change in perception as responsible for the emergence of seriation is to confuse cause with effect.

Before arguing that operation seriation cannot be accounted for by the development of perception, we ought logically to show that it is legitimate to draw a distinction between perception and cognition in the first place. The arguments usually advanced for the distinction take the form of showing that two related phenomena follow opposite developmental courses, it being argued that one of them is 'perceptual' and the other one more dependent on central 'cognition'. Thus Pollack (1969) distinguishes between Type I visual illusions, which show no relation to intelligence and diminish during childhood, and Type II illusions, which are correlated with intelligence and increase with age. Similarly, Inhelder (1962) discusses the case of the child's comparison of the lengths of two parallel lines. When the lines are aligned, thusly  , the task is a 'perceptual' one, and performance on it deteriorates with age. But when one of the lines is displaced relative to the other,  , the task demands more in the way of cognitive capacity and understanding of the notion of length, and performance improves with age. Wohlwill (1962) has gone so far as to propose the shift from dominance-by-perception to inference-making as a fundamental dimension of cognitive development.

Inhelder and Piaget (1964) argue that since the perceptual "good form" of a seriated line depends upon the sensory-motor scheme of arranging the rods in order, the good form itself can contribute little to the ability to seriate. Firstly, seriation itself changes the child's perception: "The internalisation of ordering activities transforms the child's perception (and representation) of the series because he attributes to that perception the results of his own mental activity" (Elkind, 1964). Secondly, the order of the series cannot help the child if he cannot perceive it. As Furth puts it,

"When a four-, a six-, and an eight-year-old look at an ordered series made with sticks of decreasing length, the underlying respective structures of intelligent behavior explain to a great extent what these children psychologically perceive. ... The ordered arrangement of the sticks, even though present objectively ... does not exist for the four-year-old who fails to perceive an ordered arrangement different from an unordered arrangement. The four-year-old's intelligence cannot assimilate seriated order."

(Furth, 1969, pp. 13-14)

Inhelder and Piaget bring various empirical arguments to bear on the point. Using series of rods with monotonically increasing or decreasing differences, so that the profile of the series is parabolic, they show that not until the age of 9 or 10 can children use the global shape of the series to decide which of two distinct pairs of adjacent rods differ by the greater amount. Relevant also is their well-known finding that at an appropriate age a child's memory for the appearance of a seriated line can actually improve with time, recall being better after six months than after just one week (Piaget, Inhelder & Sinclair, 1968; reported in Inhelder, 1969; Piaget & Inhelder, 1969). More directly, they show that before the age when he can construct the series for himself, the young child cannot draw a seriated line correctly even when a model is present in front of him (Inhelder & Piaget, 1964, Chapter 9).

As with perception, there is a similar temptation to treat language development as being responsible for seriation. Kendler and Kendler (1967), for example, point out that language provides the human animal with a ready-made mediational system, one that is immediately available for each new task. But they are surely mistaken in attributing success on the more advanced tasks, such as seriation, to the existence and use of language. Piaget has shown again and again that operationality does not depend on language; but that on the contrary, the concrete operations supply the cognitive structures needed for the child to use language correctly. The most direct study of the relationship between language and operational thinking has been carried out by Sinclair (1967, 1969). Studying the problems of conservation and seriation, she found that certain higher-order linguistic devices were used in a situation only by children who succeeded on the task. For example, in the conservation task, successful children tended to describe the containers in terms of comparatives like 'bigger', 'taller', whereas the younger children used only the absolutes 'big' and 'tall'. It could not have been the case that the more advanced language was responsible for the superior performance, since, firstly, training the children to use the appropriate verbal expressions did not lead to operational progress. Secondly, the pre-operational failures were not due to linguistic ignorance, because all the children in the study could both comprehend comparatives and also use them spontaneously in situations which they could master intellectually. Piaget's position also receives support from the research into the thinking of deaf children (reviewed by Furth, 1964, 1966, 1971), which indicates that operational thinking is not contingent upon the possession of language.

Piaget and information processing

Since what we are doing in this paper is basically to carry out an information processing analysis of a Piagetian task, it becomes relevant to examine briefly the relationship between Piagetian theory and the information processing approach. Certain broad similarities between the two disciplines have been noted before, by Hunt (1961), Simon (1962), Flavell & Wohlwill (1969), and Cellier (1972), among others, and we will not duplicate here what they have to say. Rather we will compare the respective approaches on two issues: the role of the environment, and the nature of symbolic representation.

In Piaget's view of development, a child's mental structures grow by interiorising³ the observed effects of his own actions on the environment; and this aspect of cognitive growth will have to play a key part in any future information processing theory of development. Conversely, the information processing use of a 'task analysis' to discover the demands of the task on the organism that performs it, and the extent to which the organism's behaviour is determined by these demands, enables us to understand rather more clearly certain issues raised in the developmental literature. To take two examples:

(1) Flavell (1963) refers to the logical groupements as capturing the objective properties of the particular concrete operation each represents. Thus the six rules of Grouping V express the properties of order as a transitive, asymmetric relation. But to the extent that this is so, it follows

³ For the difference between 'interiorise', and 'internalise', see Furth (1968, 1969).

that the groupements describe the environment and cannot tell us anything about the child other than that he has adapted to that aspect of the real world. (Compare the discussion of adaptation in Simon, 1969.)

(2) One of the categories of diagnostic error that Smedslund (1963, 1969) warns against is that the child may perform the task by "irrelevant procedures", that is, by employing a process other than the one the experimenter is interested in. To avoid this error, the experimenter has to understand his task so thoroughly that he is familiar with all the different possible ways of performing it and can recognise the symptoms of the "irrelevant" ones. In other words, he has to carry out a task analysis! The need for such an analysis, and the consequences of doing without it, are vividly illustrated by the controversy between Braine and Smedslund over the age at which transitivity first appears (Braine, 1959, 1964; Smedslund, 1963, 1965). It happened that the task chosen by Braine could be solved by processes other than inferential transitivity, as Smedslund pointed out. But neither protagonist ever considered the whole range of conceivable solution methods, and it was left for Murray and Youniss (1968; Youniss & Murray, 1970) to make a further change in the task presentation in order to eliminate another "irrelevant procedure" that both Braine and Smedslund had overlooked.

It seems appropriate here to introduce a note of cautious pessimism. If one can accept Furth's (1968, 1969) analysis at face value, there may exist a wide and possibly unbridgeable gap between the two disciplines. One of the difficulties Furth notes is that a grasp of Piaget's theory is made difficult for English-speaking readers if only because the different patterns of meaning of words in English and French encourage one to ignore distinctions

that are crucial for an adequate understanding of Piaget's system, or indeed for any understanding that avoids seeming to reduce the system to a "dispensable artifact" (Furth, 1969, p. 92). According to Furth the Piagetian system is at odds with the "majority of psychologists of different persuasions [who] base the developmental and theoretical explanation of mature thinking processes on the presence of symbols" (Furth, 1967). That description certainly includes most information processing psychologists, and indeed the approach taken by Newell & Simon (1971) could quite accurately be re-labelled as a 'symbol processing' theory. For Piaget, the internal symbolic representation is something that is constructed and comprehended by a distinct, active knowing structure that is both logically and chronologically prior. "All symbols refer directly to some object of knowing and only via the knowing to the external event"--and this is the direct opposite of the view that "the external world is known through symbols" (Furth, 1969, p. 92). The relation between the knower, on the one hand, and the symbol structure or image on the other, is similar to that between the knower and the external world (which is what makes introspection possible). The symbols derive their representational power--their meaning--only through the knowing operation which produces them in the first place.

To an information processing psychologist such an argument must at first glance look like a case of the fallacy-of-the-homunculus run riot: but that may simply be a consequence of the kind of misunderstanding Furth discusses. Certainly, even for computer simulation, there are attractions in Piaget's unfamiliar standpoint. Whereas for Newell & Simon (1971) thinking is symbol manipulation, for Piaget symbol manipulation is just one of several modes of thinking, and is mediated by a non-symbolic mechanism. Such a scheme offers attractions for modelling behaviour on sensory-motor or even concrete reasoning tasks, especially since the approach adopted by Newell & Simon makes it difficult if not impossible to deal satisfactorily with the non-symbolic aspects of cognitive behaviour.

At any rate, the present writer finds himself in disagreement with some earlier discussions which maintain that implementing Piaget's theory of cognitive development in the form of a computer program poses no problems that we do not know how to solve. Frijda's list of the remaining problems includes: "How to simulate the processes of abstraction and equilibration; how to make the right operations and concepts available to the program at the right time; and how to formalise the child's environment" (Frijda, reported by Gascon, 1969; quoted from Cell  rier, 1972). At least these latter two would appear to be currently outstanding problems in Artificial Intelligence research. He may be correct in his view that they are only "practical" problems, not "problems of principle", but it is difficult to see how that makes them any the nearer to being solved. We assert that at this time it remains unclear how to 'translate' Piaget into the language of information processing, and that it is far from obvious just what a "computerised Piaget" would look like (see Quillian, Wortman, & Baylor, 1965; Gascon, 1969).

In their chapter on Flexibility in Hindsight and Foresight, in which they come perhaps as close as they ever do to actually discussing the processes underlying cognitive developmental change, Inhelder & Piaget (1964) end with the sentence "We can now see the causal mechanisms involved". If only we could agree with them!

3. A Detailed Example: S315 on the Seriation Task

As an illustration of the various ideas discussed so far, this part of the paper presents an analysis of subject S315, a five year old boy who was video-taped while working on the seriation problem. We first describe his problem-solving behaviour at three levels of detail, including with each level any relevant general comments. We then proceed to analyse the behaviour and write a PS to model it.

The videotape was made in 1969 at the University of Pittsburgh's Learning Research and Development Center for part of the evaluation of a teacher-training program. The instructions given to the subject are far from ideal, but this need not bother us since he clearly understands the task. The rods were of uniformly increasing length, and were arranged initially lying side by side in the order F, D, E, B, A, C--A being the longest rod and F the shortest.

Level 1 protocol

S315 takes 21 seconds to solve the seriation problem, repeatedly picking the shortest rod remaining in the Pool and standing it at the right hand end of the growing Line. Twice he makes an error in selection, picking the next-to-smallest rod instead of the smallest, but both times he notices his error as he is placing the rod, returns it to the Pool, and picks the correct one.

There is little to say about this level of description except to warn the reader against drawing unjustified inferences from it. In Piagetian terms, S315 appears to be transitional between stages II and III of empirical seriation (Piaget, 1952). He tries to use the operational method of selecting always the smallest rod, but without quite having mastered the technique.

Level 2 protocol

Teacher: "See these? Can you put them in a line, from the smallest to the biggest? Can you start with the smallest and make a line ... stand them up?" (During the second sentence she twice strokes the table from left to right where the Line will be.)

S315 picks up rod F--initially the leftmost--with his left hand, and moves it to its position, still horizontal. He then reaches back and stands it on end. (Teacher: "That's right!") He picks up rod E with his left hand and places it next to rod F with his right. Similarly with D. He then picks up B in his left hand and with his right moves it to near rod D but not in the Line; without having let go, he transfers it back to his left hand and returns it to its original position. He then picks up C in his right hand and places it. He picks up A and places it, but without moving his hand away he picks it up again and returns it to the Pool. He picks up B with his left hand and places it with his right, meanwhile using his left hand to stop A which has been rolling towards the edge of the table. Finally he picks up and places A.

Teacher: "Oh, that's very good ..."

Two points are worth noting about this protocol. First, S315 places all rods after F with his right hand, after transferring them if necessary from his left; presumably because this is less clumsy than placing them with his left hand when building the Line towards the right.

The second is that he slaves himself to the details of the teacher's instructions. He constructs the Line "from the smallest to the biggest" rather than in the opposite order. He builds it from the smallest at the left to the largest at the right, the same direction as the teacher strokes the table when she is giving the instructions. And he builds it at just the place on the table that she indicates, even though he has difficulty reaching so far and when he leans forward has to support himself on the table with his free hand.

Level 3 protocol

The level 3 description is based on a detailed transcript of the videotape which includes timings. Rather than presenting it here in its entirety, we simply extract the part describing the two errors:

... He seems deliberately to select B. He grasps B with his left hand and transfers it to the right. He places it near the Line, seems to pay attention to it, then--without having let go--he picks it up, transfers it to his left hand, and returns it to the Pool. His right hand picks up and places C, apparently with great certainty. With his right hand he picks up A--again the selection is apparently deliberate--and places it. He lets go of the rod but does not move his hand away. He is definitely looking at the Line. He picks up A again and returns it ...

Several of the terms occurring in this protocol--"deliberately", "certainty", and so on--are more interpretative than objective, and their appearance there requires some justification. We use them mainly because we are neither willing nor able to transcribe from the videotape the many subtle cues that contribute to an impression of certainty and so on. Partly, we use them because they provide an appropriately detailed description of S315's behaviour. At level 3 we are not interested in his eye- and head-movements themselves but in what he pays attention to. We also suppose that, if put to the test, a panel of independent observers would largely agree on these interpretations--but that supposition remains an item of faith rather than a known fact.

The seriation strategy

Before writing a PS to model S315's seriation behaviour we have to decide what strategy we think he is using. At first glance it seems that he is attempting always to pick up the smallest rod remaining in the Pool and then place it, as was suggested by the level 1 protocol. However, a careful examination of the selection errors shows that conjecture to be wrong.

The argument hinges on two pertinent features of the errors. We note firstly that both errors occur towards the end of the seriation, when there are only three and then two rods remaining, respectively, and therefore involve only the longer rods (A, B, and C). Secondly, S315 notices his mistakes when he is comparing the wrongly selected rod against the part of the Line already constructed; in particular, he does not notice them while looking at the Pool. Now if he were in fact always trying to pick the smallest rod remaining, although it is quite plausible that he would commit some errors, it is hard to understand why he should make a mistake when there are only two rods left. When he makes his second error rods A and B are lying side by side, and he appears to pick A carefully rather than impulsively. If he were deliberately trying to pick the smallest rod this should be the easiest choice of all, and therefore the least likely to go wrong. Moreover, the place for S315 to check the rod and exchange it if necessary would be at the Pool, right after making his selection: he should at that point verify that the rod in his hand is indeed the smallest. But what he does is quite different. The incorrect rods are checked and rejected at the Line: he compares both B and A to the rods already placed without even moving his hand away.

These considerations suggest an alternative strategy. Rather than intentionally selecting the smallest rod, we think that S315 is making use of the uniform difference in lengths to estimate from the part of the Line already built how tall the next rod should be. He then goes to the Pool with that length in mind and picks a rod of that particular size. This strategy differs from the 'select.smallest' strategy mainly in where it assumes the information governing the next selection is located. From this alternative

technique S315's verification procedure follows naturally, since it is at the Line that he has the necessary information to check his selection. The errors also become understandable--for two reasons. First, the equal increments mean that the longer the rods, the less discriminable they are, so that to confuse the largest two rods is now the most (instead of the least) likely way to make an error. Second, the task is considerably more demanding than a simple Select.smallest. S315 has to estimate the size of the rod by perceptual extrapolation from the Line, retain this estimate while transferring his attention and gaze to the Pool, and then try to match it against one of the remaining rods.

In the final section of this paper we will discuss ways to distinguish experimentally between these two proposed strategies. For the time being, however, for purposes of analysis we assume that S315 is indeed performing this estimation procedure.

Production system for S315

In the analysis which follows we will be concerned only with the repetitive, cyclic phase of the seriation behaviour, which involves selecting rods from the Pool one by one and placing them in the Line. We will make no attempt to model accurately--or, for that matter, at all--the beginning and end of the seriation process. The reason for this restriction is that our one sample of S315's behaviour simply does not provide us with enough information to distinguish among the several conceivable ways in which he might be starting and finishing.

The protocols make clear that S315 cannot be using the same goal structure as our idealised PS. If he were alternately setting up the goal of getting the next rod and then, once he had it, the goal of placing it (as does

our hypothetical child), his behaviour on detecting an error would be quite different. Once he realised that he had an incorrect rod, since his immediate goal would be to place it anyway he would probably guess its correct position and put it there, leaving a gap to be filled by the "next" rod.

(Adults can be observed to do precisely this under analogous circumstances.)

That he goes to all the trouble of returning the rod exactly where he found it before looking for the correct one, indicates that his behaviour must be driven by a need to place the rod next in the Line:

(S1) Place.all; => Set.goal [Place]

If the goal is to place all the rods, i.e. to do the seriation task, then set up the goal of placing one of them.

Presumably he goes and gets a rod because he cannot place one unless he has one:

(S2) Place; Hand = 0 => Set.goal [Get[]]

If the goal is to place a rod and there is not one in the Hand, then set up the goal of getting one.

We are assuming that he decides which rod to pick next by extrapolating from the Line:

(S3) Get[] => Estimate

If the goal is to get some unspecified rod, then estimate from the rods already in the Line the height of the next one. Use this information to specify the goal further.

Once he has this further information, S315 may be able actually to pick the next rod:

(S4) Get[particular]; At[Pool] => Select[particular]

If the goal is to get a rod of some specified size, and if the Hand is at the Pool, then pick up the rod of that size.

Since we are not interested in the details of the complex operation of placing the rod, checking the rod, and possibly rejecting it, we simply write the rules as follows without intending to commit ourselves seriously to their details:

(S5) Place; Hand = 1; At[Line] => Compare

If the goal is to place a rod and there is one in the Hand, then if the Hand is at the Line compare the rod with the ones already placed.

(S5') Place; Compare + + => Insert

If the goal is to place a rod, and if the new rod has just been compared against the Line and found to extend it correctly (i.e. the new "step size" is the same as the previous ones), then insert the rod in the Line.

(S5'') Place; Compare + - => Reject

Similarly, if the new rod has just been found to be incorrect, return it whence it came.

Finally, we need two instances of the general rule we discussed earlier to move the Hand where it needs to be if it is not already there:

(S6) Get[particular] => To[Pool]

(S7) Place; Hand = 1 => To[Line]

If the goal is to get a rod of specified size (and if the Hand is not at the Pool), then move the Hand to the Pool. Similarly, if the goal is to place a rod already in the Hand (but the Hand is not at the Line), move it to the Line.

The rules are normally invoked in the sequence S1-S2-S3-S6-S4-S7-S5-S5'- ... but when an error occurs S5' is followed immediately by S2. The goal of Get is satisfied by Select, and the goal of Place by Insert.

The nature of the PS may become clearer if we see how the various rules can be grouped in various ways. (1) Since the overt behaviour consists of alternately Picking up a rod and then placing it, we can view rules S3, S4, and S6 as being responsible for picking up the rod, and S2, S5, S5', and S7 for placing it. S1 is a kind of "pacemaker" whose job is to initiate each cycle of activity. (2) Alternatively, the PRs can be grouped by their function. S1 and S2 are for setting the goals, S3 embodies the selection strategy employed, S4, S5, and S5' actually manipulate the rods, and S6 and S7 are for moving from place to place. (3) A third view is that some of the rules are required in order to make it possible for other rules to be invoked. For example, S5' is the rule that actually builds the seriated line, but before it can operate the Hand must be at the Line, holding the appropriate rod. So S7 ensures that the Hand is at the Line at the right time, and S2 is responsible for seeing that it holds a rod. In a similar way, S3 and S6 serve to ensure that the pre-conditions of S4 are satisfied.

Discussion of the production system

It is revealing to compare the present PS with the idealised one given in section 1 of this paper and notice the enhanced psychological--or more precisely, developmental--face validity of this real one. Each of its PRs possesses the kind of local plausibility that we discussed as a desirable attribute of PSs. Rule P1 can be seen as a specialised version of a general rule the child has for dealing with sets of objects one at a time, quite

independent of this particular task: something like "If you want to do something to each of a set of objects, do it to one of them." The rule keeps applying until there are no objects left. S2 simply states that if one ever wants to do something that requires a rod, then a good idea is to have a rod to do it with. S3 indicates that if one needs to know what size rod to pick next, one way to find out is to examine the rods already in the Line. S4 expresses a standard means for getting hold of a rod of particular size: Namely, to select one of that size when the opportunity arises. And S5 and S5' provide a similar means for achieving the goal of placing a rod correctly. (S6 and S7 should be old friends by now.)

This property of local plausibility can be turned around and used as a criterion by which to evaluate a proposed PR or PS, and in this guise it provides further grounds for supposing that S315 is indeed employing the strategy of estimation, rather than that of selecting the smallest rod. To be doing the latter means that he would have some PR of the form

(S3') Get[] => Get[smallest]

If the goal is to get the next rod, then specify
it to be the goal of getting the smallest rod

and it is difficult to see where such a rule could have come from. "Smallest" appears to bear little intrinsic relation to "next-ness".

Nevertheless, S315's PS supplies an insight into how the "operational" strategy of selecting the smallest rod could develop. The key point to notice is that the two strategies generate virtually identical overt behaviour: in both cases S315 can be described as always picking the smallest rod from the Pool and then placing it. This means that a young child who is following the

a priori plausible rules that comprise the estimation strategy does in fact pick the smallest rod from the Pool, even though that is not what he is "doing". He could thus learn to do a 'Select.smallest' by a process of "reflective abstraction" from his own behaviour (Piaget, 1970). The developmental sequence suggested here is that by following one set of rules, the child is led to generate behaviour more economically described by a different set of rules; these rules he in turn incorporates, as an induction from his own actions. In going through the laborious and difficult business of estimating the size of the next rod, he is repeatedly led to pick the smallest one, so that eventually he comes to do just that.

The difference between the two sets of rules is that whereas the first possesses local plausibility that allows the individual PRs to be "assembled" from many different and unrelated origins, the second implicitly embodies certain logico-physical properties of the external world. The new rules develop not through the child's observation or deduction of these properties, but from the regularities in his overt behaviour. Thus this protocol provides material for a speculative account of how, in terms of a PS, a child might undergo the Piagetian process of growing new cognitive structures through the interplay between his existing abilities and the properties of his environment. The next task is to see how one sets about investigating a child's production system experimentally.

4. Validation of Production Systems

Once a production system (PS) has been written to account for certain features of a segment of behaviour, the question arises as to how to confirm the PS against independent evidence. In this section we discuss two ways of utilising further information in order to choose between alternative PSs. So far as we know, this kind of validation has never yet been used. The only evaluation procedure mentioned by Newell & Simon (1971) is limited to measuring how thoroughly a PS accounts for the behaviour it purports to describe (Chapter 6).

Validation techniques

As was pointed out in Section 1, a PS by itself is incapable of actually predicting behaviour; to do that, the properties of the environment must also be specified. Since this means that the same PS can generate different behaviour in different situations, it follows that a PS written solely to mimic a given sample of a child's behaviour can to some extent also be used to predict his behaviour in a different environment. This prediction can then be used as a basis for evaluating the hypothesised PS or for distinguishing between rival ones.

To continue with the example of subject S315 that we have just been examining, we claimed there that whenever he picked a fresh rod, he was not in fact deliberately selecting the smallest one remaining but rather was looking for a rod of a particular estimated size. The PS was written accordingly. Suppose now that we were to set him to work on a similar seriation task, but this time with one of the middle-size rods, D say, removed from the

set. What would he do? If he were really simply picking the smallest rod remaining, he should go ahead and select C right after E without any trouble; for each of these in turn would indeed be the smallest rod left in the pool. On the other hand, if he were actually estimating the size of the next rod--as we have claimed--then after he places rod E we can expect him to be very unhappy about not being able to find the rod he wants next. Note that in this latter case the PS does not actually specify what S315 will do; it merely predicts that the Estimation strategy will lead to an evident disruption of his ongoing behaviour, whereas the Select.smallest strategy will not. By virtue of this differential prediction, the task of seriation-with-one-missing provides an example of how a subsidiary experiment can help decide between alternative strategies that the child might be following.

The idea of probing a child's understanding by having him perform some task in circumstances that differ from his expectations is hardly original. Many examples can be found in the developmental literature. Perhaps the best known is Smedslund's (1961) use of sleight-of-hand in testing how firmly children are convinced of the conservation of weight. In the usual context of asking whether the weight of a ball of clay is changed when it is rolled into a sausage, he had the children compare two initially equal lumps on a beam balance. By deftly removing a small piece of clay while it was being rolled, he managed to produce an apparent violation of conservation, and the child was categorised as a conserver or non-conserver according to his response to this violation. In a similar fashion, he surreptitiously changed the comparison stick in the process of a length transitivity problem (Smedslund, 1963). Charlesworth (1964, 1969) regards the child's reaction

to an environmental "surprise"--i.e. a violation of his expectations--not only as a powerful tool for investigating his cognitive structures, but also as a principle means whereby development takes place.

If a child can validly be described as having the rule

(P) C => A

then, provided that P is not 'blocked' by another PR of higher priority, action A should be taken whenever condition C is satisfied--that is what having the rule means. If we temporarily ignore the matter of conflicting rules, production P should be evoked (a) no matter what the environment in which C holds, and (b) no matter what actions the child has just taken to make C true. It may be difficult at this point to see any difference between (a) and (b), and indeed these two conditions are but the two sides of the same coin, but nonetheless emphasising each in turn leads to two distinct, complementary techniques for validation. The first approach leads to a direct confirmation of individual productions, whereas the second yields a somewhat weaker support for the PS as a whole:

(a) By concentrating on verifying that rule P is evoked regardless of the properties of the environment, we are led to the idea of intervening in the task in order to change the environment unexpectedly. The critical test then consists of showing that P is evoked when C is made true. To take a simple example, the PS for S315 included the rule

(S7) Place; Hand = 1 => To[Line].

Normally the left hand side becomes true when the child chooses and then grasps a suitable rod; but if instead the experimenter simply hands him a rod at an appropriate moment, the rule should be obeyed immediately. In

this way we arrive at a sort of validation-in-depth technique: intervene in the task unexpectedly, and verify that the appropriate rule is obeyed.

(b) On the other hand, by concentrating on verifying that rule P is evoked no matter what the child did to make C true, we are led to the idea of posing a number of different but related tasks in which quite different actions by the child result in condition C being satisfied. Again, the critical test consists of seeing whether P is the next production to be invoked. By way of illustration, there exist any number of block manipulation tasks in which the subject is required to place blocks one by one to meet certain given constraints; and in all of them it is the case that if the immediate goal is to place the next block, the appropriate action is to go off and find it. We can express this by the production

Place; Hand = 0 => Set.goal[Get[constrained]]

and an instantiation of this rule should appear in the PSs for each of the separate tasks. So by following this line of reasoning we arrive at a sort of validation-in-breadth technique: write separate production systems for the child's behaviour on a number of related tasks, and verify that corresponding functions are performed by the same individual rules in each.

In order to realise these twin approaches, we have devised a whole collection of subsidiary tasks suitable for investigating a wide range of seriation abilities. The tasks divide themselves into loosely-defined categories: (1) Many are variants on the basic task of length seriation, using different numbers, sizes, and arrangements of blocks. (2) There are seriations along dimensions other than length: diameter (of wooden discs), numerosity, and weight. (3) There are different 'modes' of seriation: with an extra rod,

or a missing one (as discussed earlier); with the line built behind a screen; or with the seriation performed blindfold. (4) There are other tasks using the blocks, such as copying a given pattern of coloured blocks, or constructing a line out of blocks of alternating colours. (5) There are 'probe' tasks, in which the subject has to hand the experimenter the larger of two blocks, or insert a new rod or correct an error in a given series. (6) And finally there are the 'interventions', in which the experimenter changes the child's environment, for example by unexpectedly handing him a rod to place.

These validation techniques suggest a mode of experimentation somewhat foreign to experimental psychology, though it is familiar enough to clinicians and is compatible with Piaget's interview techniques. In his first experimental session, an individual child is asked to perform a straightforward seriation task and solve a number of other simple problems, his behaviour all the while being recorded on videotape. When the experimenter carries out a preliminary analysis of the recorded behaviour, and tries to write a production system, he will probably come up with questions which cannot be answered on the basis of this first round of data. So, perhaps a week after his first session, the child is called back to face a variety of tests specially designed to yield evidence for settling the issues still to be resolved. If the experimenter requires yet more data, the child can return for a third or even a fourth session.

Thus each subject receives a different sequence of problems, chosen not randomly but in order to help clarify particular features of his PS. In this sense the investigation follows the overall plan of an adaptive experiment; but the times involved are far longer than that term normally suggests.

Proposed research

The proposed thesis research will consist of an examination of seriation behaviour in children using the methodology just outlined. The work will be exploratory, in the sense that we expect to follow up on only a few of the issues we raise, and also frankly opportunistic in the sense that any interesting clues that arise will be pursued for whatever they are worth. We will avoid completely the kind of systematic study that requires large numbers of subjects and trials together with appropriate statistical controls; so that, for instance, we will not attempt to investigate the perceptual discriminability of the seriation blocks we will be using.

The research will be carried out at the Teacher Training laboratory of the University of Pittsburgh's Learning Research and Development Center, making use of children from the Center's subject pool. The laboratory is equipped with a one-way viewing screen and two-camera video-recording equipment. A digital electronic counter has been arranged in front of one of the cameras in such a way that, when superimposed on the other picture, each TV frame carries a five-digit identifying number. This proves to be of great help in transcribing data from the videotapes and for measuring times.

It is intended to analyse in some detail about half-a-dozen children of different levels of ability: probably two children who cannot seriate, two who seriate by trial-and-error, and two who seriate operationally. The children's behaviour will be analysed in a way intended to emphasise the similarities and differences between the production systems for the different levels of performance; hopefully this will make it possible to extract from the data the greatest possible understanding of how children acquire their ability to seriate.

References

- Battisti, C., & Simmons, R. Transitivity and seriation performance for 5, 6, and 7 year old children. S.U.N.Y. at Buffalo, Child Study Center Bulletin, 1968, 4, 64-69.
- Bellugi, Ursula, & Brown, R. (Eds.) The acquisition of language. Monograph of the Society for Research in Child Development, 1964, 29, No. 92.
- Bloom, Lois. Language development: Form and function in emerging grammars. Cambridge, Mass.: M.I.T. Press, 1970.
- Braine, M. D. S. The ontogeny of certain logical operations: Piaget's formulation examined by non-verbal methods. Psychological Monographs, 1959, 73 (5), (Whole #475).
- Braine, M. D. S. Development of a grasp of transitivity of length: a reply to Smedslund. Child Development, 1964, 35, 799-810.
- Bruner, J. S., Olver, R. R., & Greenfield, P. M. Studies in Cognitive Growth. New York: Wiley, 1966.
- Cell  rier, G. Information-processing tendencies in recent experiments in cognitive learning: Theoretical implications. In S. Farnham-Diggory (Ed.), Information processing in children. New York: Academic Press, 1972.
- Charlesworth, W. R. Development and assessment of cognitive structures. Journal of Research in Science Teaching, 1964, 2, 214-219.
- Charlesworth, W. R. The role of surprise in cognitive development. In D. Elkind & J. H. Flavell (Eds.), Studies in Cognitive Development. New York: Oxford University Press, 1969.
- Chomsky, Carol. The Acquisition of Syntax in Children from 5 to 10. Cambridge, Mass.: M.I.T. Press, 1969.

- Elkind, D. Discrimination, seriation, and numeration of size and dimensional differences in young children: Piaget replication study VI. Journal of Genetic Psychology, 1964, 104, 275-296.
- Flavell, J. H. The Developmental Psychology of Jean Piaget. Princeton, N.J.: van Nostrand, 1963.
- Flavell, J. H., & Wohlwill, J. F. Formal and functional aspects of cognitive development. In D. Elkind & J. H. Flavell (Eds.), Studies in Cognitive Development. New York: Oxford University Press, 1969.
- Fodor, J. A. Psychological explanation: An introduction to the philosophy of psychology. New York: Random House, 1968.
- Frijda, N. H. Problems of computer simulation. Behavioral Science, 1967, 12, 59-67.
- Furth, H. G. Research with the deaf: Implications for language and cognition. Psychological Bulletin, 1964, 62, 145-164.
- Furth, H. G. Thinking without language: Psychological implications of deafness. New York: Free Press, 1966.
- Furth, H. G. Concerning Piaget's view on thinking and symbol formation. Child Development, 1967, 38, 819-826.
- Furth, H. G. Piaget's theory of knowledge: The nature of representation and interiorisation. Psychological Review, 1968, 75, 143-154.
- Furth, H. G. Piaget and knowledge: Theoretical foundations. Englewood Cliffs, N. J.: Prentice-Hall, 1969.
- Furth, H. G. Linguistic deficiency and thinking: Research with deaf subjects 1964-1969. Psychological Bulletin, 1971, 76, 58-72.
- Gascon, J. Modèle cybernétique d'une sériation de poids chez les enfants. Université de Montréal, Institut de Psychologie, M.C.P. #2, 1969.

- Glick, J., & Wapner, S. Development of transitivity: some findings and problems of analysis. Child Development, 1968, 39, 621-638.
- Gyr, J. W., & Fleischer, C. Computer-assisted studies of Piaget's three-stage model of cognitive development. Psychological Reports, 1967, 20, 165-166.
- Hunt, E. B. What kind of computer is man? Cognitive Psychology, 1971, 2, 57-98.
- Hunt, J. McV. Intelligence and Experience. New York: Ronald Press, 1961.
- Hunt, J. McV. Impact and limitations of the giant of developmental psychology. In D. Elkind & J. H. Flavell (Eds.), Studies in Cognitive Development. New York: Oxford University Press, 1969.
- Inhelder, B. Some aspects of Piaget's genetic approach to cognition. In W. Kessen & C. Kuhlman (Eds.), Thought in the young child. Monograph of the Society for Research in Child Development, 1962, 27 (2). (Re-issued by the University of Chicago Press, 1970).
- Inhelder, B. Memory and intelligence in the child. In D. Elkind & J. H. Flavell (Eds.), Studies in Cognitive Development. New York: Oxford University Press, 1969.
- Inhelder, B., & Piaget, J. The early growth of logic in the child: Classification and seriation. New York: Harper and Row, 1964.
- Kendler, T. S., & Kendler, H. H. Experimental analysis of inferential behavior in children. In L. P. Lipsitt & C. C. Spiker (Eds.), Advances in child development and behavior. Vol. 3. New York: Academic Press, 1967.
- Kessen, W. "Stage" and "Structure" in the study of children. In W. Kessen & C. Kuhlman (Eds.), Thought in the young child. Monograph of the Society for Research in Child Development, 1962, 27 (2). (Re-issued by the University of Chicago Press, 1970.)

- Klahr, D., & Wallace, J. G. An information processing analysis of some Piagetian experimental tasks. Cognitive Psychology, 1970, 1, 358-387.
- Klahr, D., & Wallace, J. G. Class inclusion processes. In S. Farnham-Diggory (Ed.), Information processing in children. New York: Academic Press, 1972.
- Klima, E. S., & Bellugi, Ursula. Syntactic regularities in the speech of children. In J. Lyons & R. Wales (Eds.), Psycholinguistic papers. Edinburgh: Edinburgh University Press, 1966.
- Knifong, J. D. The representation of cognitive structures of 4½ year old children. Ph.D. thesis, University of Illinois, Urbana, Ill. 1971.
- Kofsky, E. A scalogram study of classificatory development. Child Development, 1966, 37, 191-204.
- McManis, D. L. Conservation and transitivity of weight and length by normals and retardates. Developmental Psychology, 1968, 1, 373-382.
- McManis, D. L. Conservation, seriation, and transitivity performance by retarded and average individuals. American Journal of Mental Deficiency, 1970, 74, 784-791.
- McNeill, D. Developmental psycholinguistics. In F. Smith & G. A. Miller (Eds.), The genesis of language: A psycholinguistic approach. Cambridge, Mass.: M.I.T. Press, 1966.
- McNeill, D. The acquisition of language: The study of developmental psycholinguistics. New York: Harper & Row, 1970.
- Menyuk, Paula. Sentences children use. Cambridge, Mass.: M.I.T. Press, 1969.
- Miller, G. A., Galanter, E., & Pribram, K. H. Plans and the structure of behavior. New York: Holt Rinehart & Winston, 1960.

- Murray, J. P., & Youniss, J. Achievement of inferential transitivity and its relation to serial ordering. Child Development, 1968, 39, 1260-1268.
- Newell, A. On the analysis of human problem solving protocols. In Calcul et formalisation dans les sciences de l'homme. Proceedings of an international symposium held in Rome, July 1966. Paris: Éditions du Centre National de la Recherche Scientifique, 1968.
- Newell, A., Shaw, J. C., & Simon, H. A. Elements of a theory of human problem-solving. Psychological Review, 1958, 65, 151-166.
- Newell, A., & Simon, H. A. An information processing theory of human problem solving. Englewood Cliffs, N. J.: Prentice-Hall, 1971.
- Osgood, C. E. Method and theory in experimental psychology. New York: Oxford University Press, 1953.
- Piaget, J. The child's conception of number. New York: Humanities Press, 1952.
- Piaget, J. The construction of reality in the child. New York: Basic Books, 1954.
- Piaget, J. Genetic epistemology. New York: Columbia University Press, 1970.
- Piaget, J., Inhelder, B., & Sinclair, H. Mémoire et intelligence. Paris: Presses Universitaires de France, 1968.
- Piaget, J., & Inhelder, B. The psychology of the child. New York: Basic Books, 1969.
- Pollack, R. H. Some implications of ontogenetic changes in perception. In D. Elkind & J. H. Flavell (Eds.), Studies in Cognitive Development. New York: Oxford University Press, 1969.
- Proscura, E. V. The role of teaching in the formation of seriation actions in preschool children. Voprosy Psichologii, 1969, 15, 37-45. Cited by I. Z. Holowinsky, Seriation actions in preschool children. Journal of Learning Disabilities, 1970, 3, 457-458.

- Quillian, M. R., Wortman, P. M., & Baylor, G. W. The programmable Piaget: Behavior from the standpoint of a radical computerist. Carnegie-Mellon University, Department of Psychology. Unpublished MS, 1965.
- Shantz, C. U. Effects of redundant and irrelevant information on children's seriation ability. Journal of Experimental Child Psychology, 1967, 5, 208-222.
- Simon, H. A. An information-processing theory of intellectual development. In W. Kessen & C. Kuhlman (Eds.), Thought in the young child. Monographs of the Society for Research in Child Development, 1962, 27 (2). (Re-issued by the University of Chicago Press, 1970.)
- Simon, H. A. The sciences of the artificial. Cambridge, Mass.: M.I.T. Press, 1969.
- Sinclair, H. Acquisition du langage et développement de la pensée. Paris: Dunod, 1967.
- Sinclair, H. Developmental psycholinguistics. In D. Elkind & J. H. Flavell (Eds.), Studies in Cognitive Development. New York: Oxford University Press, 1969.
- Smedslund, J. The acquisition of conservation of substance and weight in children: III. Extinction of conservation of weight acquired "normally" and by means of empirical controls on a balance. Scandinavian Journal of Psychology, 1961, 2, 85-87.
- Smedslund, J. Development of concrete transitivity of length in children. Child Development, 1963, 34, 389-405.
- Smedslund, J. Concrete reasoning: A study of intellectual development. Monographs of the Society for Research in Child Development, 1964, 29 (2), (No. 93), 1-39.

- Smedslund, J. The development of transitivity of length: A comment on Braine's reply. Child Development, 1965, 36, 577-580.
- Smedslund, J. Psychological diagnostics. Psychological Bulletin, 1969, 71, 237-248.
- Waldinger, R. J., & Lee, R. C. T. PROW: A step toward automatic program writing. In D. E. Walker & L. M. Norton (Eds.), Proceedings of the First International Joint Conference on Artificial Intelligence. Washington, D. C., May 1969.
- Waterman, D. A. Generalization learning techniques for automating the learning of heuristics. Artificial Intelligence, 1970, 1, 121-170.
- Wohlwill, J. F. A study of the development of the number concept by scalogram analysis. Journal of Genetic Psychology, 1960, 97, 345-377.
- Wohlwill, J. F. From perception to inference: A dimension of cognitive development. In W. Kessen & C. Kuhlman (Eds.), Thought in the young child. Monographs of the Society for Research in Child Development, 1962, 27 (2). (Re-issued by the University of Chicago Press, 1970.)
- Youniss, J., & Murray, J. P. Transitive inference with nontransitive solutions controlled. Developmental Psychology, 1970, 2, 169-175.