

MODELING COGNITIVE DEVELOPMENT ON THE BALANCE SCALE TASK

Stephanie Sage

Pat Langley

The Robotics Institute

Carnegie-Mellon University

Pittsburgh, Pennsylvania 15213 USA

ABSTRACT

In this paper we describe a production system model of children's development on the balance scale task. Starting with a set of rules that makes random predictions, the system learns from its errors and improves as it gains experience. The transition mechanism is a discrimination process that searches for differences between cases in which correct predictions are made and cases in which errors are made. The stages through which the system progresses are very similar to those observed in children, so the model provides an explanation of the observed developmental trends. Since the system has no notion of torque, it never acquires the ability to completely predict the balance scale's behavior; however, it is able to learn heuristically useful rules despite its incomplete representation of the environment, much as children do.

INTRODUCTION

One of the most challenging areas facing Cognitive Science is that of cognitive development. Although the child enters the world in a nearly helpless state, within a decade he can manipulate objects, reason abstractly, and communicate with others, and he has acquired a host of other skills too numerous to mention. If we ever hope to understand the nature (and nurture) of the human information processing system, the developmental processes that lead to these diverse abilities must be understood. For a long period, research on cognitive development was almost entirely experimental, but in recent years efforts have been made to explain developmental trends in computational terms.

Production systems have become a popular framework for modeling behavior at different stages, since their modularity allows the statement of successive models that differ by only one or two condition-action rules. For example, Baylor, Gascon, Lemoyne, and Pother [1] have constructed production system models of children at various stages on Piaget's weight seriation task, and Young [2] has devised similar models for the related length seriation task. However, this work has focused on modeling behavior at a given stage, rather than explaining the transition between stages. In this paper we present a process model of development in one domain – the balance scale task. Below we describe the task, along with some earlier work in the area. After this, we discuss our model of the transition process and its implications.

THE BALANCE SCALE TASK

The balance scale task is commonly employed in the study of cognitive development, having first been used by Piaget [3]. In this task, the child is presented with a two-arm balance, having several pegs spaced evenly along each arm. Small disks of equal weight are placed on the pegs (only one peg on each side

has weights), and the child is asked to predict the direction in which the scale will move when released. The standard method for correctly making this prediction involves the notion of torque. The number of weights on a peg is multiplied by that peg's distance from the center. If one side has the greater product, that side will go down; if the products are equal, the scale will remain balanced. Lacking knowledge of this rule, children and even some adults have difficulty in making the correct predictions. However, it is clear that performance improves with age, and this trend requires an explanation.

Klahr and Siegler [4] have studied children's behavior in this domain. They found evidence for the existence of four basic stages, and successfully modeled each stage as a simple production system. For example, children in Stage 1 focused only on weight information in making predictions. The authors explained this behavior with a simple two-rule model. The first rule applied if one side had more weights, and predicted that side would go down. The second rule predicted that the scale would balance if the weights were the same. Later stages were modeled by the addition of new rules for making predictions. The second stage model included a rule that focused on distances if the weights were equal; in such cases, it predicted that the side with greater distances would go down. The Stage 3 model included a rule for selecting a side randomly when weight and distance cues were in conflict, as well a rule for dealing with cases in which the cues agreed. The model of the final stage incorporated rules for computing torque and basing the decision on the result of this calculation.

Several other studies have examined children's behavior in this domain as well. However, we have chosen to focus on Klahr and Siegler's work. Although their models were very simple, they accounted for much of the variance observed in children's behavior on the balance scale task. Moreover, the model of each successive stage differed from the previous one by the inclusion of only one or two new rules. Thus, although the authors proposed no theory of the developmental process itself, their model laid the necessary groundwork for an analysis in terms of learning mechanisms. Now let us turn to a revised stage model that takes us another step closer to an explanation of the transition process.

Figure 1 summarizes our revised model of successive stages on the balance scale task. The condition-action rules are paraphrased in English for the sake of clarity; the term *side* represents a variable that can match against either the left or right side of the scale. In addition to the first three stages modeled by Klahr and Siegler, we have included an initial random stage. Taken together, the first pair of productions (BALANCE-1 and DOWN-1) randomly predict that one of the two sides will go down, or that the sides are balanced. Klahr and Siegler found no evidence for this stage, presumably because

their subjects had moved beyond the random strategy at an earlier age. When the second pair of rules (BALANCE-2 and DOWN-2) is added to the first pair, the resulting system behaves exactly as Klahr and Siegler's model of Stage 1, provided the new rules take precedence over the first pair. Upon adding the third pair of productions (BALANCE-3 and DOWN-3), we have a model of Stage 2 behavior, provided that both of these rules take precedence over BALANCE-2. Finally, when the productions DOWN-4 and DOWN-5 are inserted (and mask DOWN-2), we have a model of Klahr and Siegler's third stage.

The revised model is in many respects very similar to Klahr and Siegler's model, but there is one major difference: *rules occurring in later stages are always discriminant versions of rules that have occurred in an earlier stage*. In other words, for every rule *s*, there exists some rule *g* in a previous stage with the same actions as *s* and a subset of *s*'s conditions. While this feature was true of some of Klahr and Siegler's rules, it was by no means true of them all. This characteristic of the revised stage model suggests a mechanism to account for the transition between these stages, which we discuss in the following section.

A MODEL OF THE TRANSITION PROCESS

The increasing specificity of the rules in successive stages suggests that they might be learned through a process of *discrimination*. Such a learning mechanism would be called when the performance system made an error, in an attempt to determine the conditions under which a prediction should be made. Potentially useful conditions can be found by comparing the situation in which a rule applied incorrectly to the last situation in which that same rule applied correctly. If some difference is found between the good and bad instances of the rule, then a variant of the rule is created that contains that difference as one or more new conditions. Because of its new conditions, the variant rule will fail to match against the

undesired case, but will still match in the desired situation. Similar discrimination learning schemes have been explored by other researchers [5, 6], so our contribution lies not so much in developing this approach to learning, as in showing how this approach can be used to explain developmental trends on the balance scale task. The details of the discrimination method have been described at greater length by Langley [7].

We have constructed a discrimination-based model of development on the balance scale task that acquires the rules presented in Figure 1. The model does not consist of the discrimination mechanism alone. It is also given the rules BALANCE-1 and DOWN-1, which provide the initial behavior upon which learning is based. In addition, the system contains a rule for comparing the sides of the balance scale on dimensions like weight and distance so that discrimination can discover conditions referring to their relative values. Finally, the model includes one rule for storing credit with the responsible rule when a correct prediction is made, and another similar production for evoking the discrimination process when an incorrect prediction is made.

For the system to improve on the balance scale task, the rules acquired in Stage 1 must be preferred over the initial random rules, and the Stage 2 and Stage 3 rules must be preferred over the Stage 1 rules. To achieve this, a rule is weakened whenever it makes an incorrect prediction, and a variant is strengthened whenever it is relearned. When they are first created, variants are weaker than their parent rules, and they must be learned a number of times before they begin to alter behavior. This approach has worked well on the balance scale task, and appears to be a generally useful technique for directing search through the space of possible rules.




Figure 1. Incremental model of behavior on the balance scale task.

The program was presented with problems selected randomly from seven basic problem types. These included problems in which only the weights differed, in which only the distances differed, in which both weights and distances were equal, in which the two cues agreed, and three types of problems in which weight and distance cues conflicted. Figure 2 summarizes the model's errors as a function of time (in units of 10 trials). Since the system begins with the two random rules BALANCE-1 and DOWN-1 and there are three basic predictions from which to choose (left down, balance, and right down), one would expect about 33 percent of the initial predictions to be correct, and this is approximately what we find. By trial 100, the system has learned (and sufficiently strengthened) the Stage 2 and Stage 3 rules, so that it makes correct predictions on all but the three conflict problems, giving a success rate of approximately 60 percent. In the case of conflict problems, the model's representation of the environment (consisting only of information about relative weights and distances) is incapable of even stating the torque rule that would correctly predict results in a consistent manner. In other words, the program's representation of the problem is inherently incomplete. However, the discrimination process is sufficiently robust to learn useful rules despite this limitation, and the system arrives at a set of rules that make correct predictions much of the time, just as children do before they are taught the torque rule.

This brings an important feature of discrimination learning to light: it allows one to learn useful rules *even if one's representation is ultimately inadequate*. Since our system has no notion of torque, it can never fully understand the balance scale task, yet it does learn rules that lead to correct predictions in many cases. Since one can never guarantee that a representation is optimal, this is a powerful feature that would be advantageous to any learning system. In addition to being interesting from an AI perspective, the ability to learn useful rules despite incomplete representations is a prerequisite for modeling human behavior on the current task, since children apparently do just that.




Figure 2. Learning curve for the balance scale task.

Although the model can never master the conflict problems, its behavior in these cases is revealing. When an error is made on a conflict problem, the system may construct variant rules, but it will be unable to discover any conditions which will make consistently correct predictions. Thus, these rules will continually be weakened and then strengthened, and the system's preference on conflict problems will oscillate between them. This effect is very similar to behavior that Klahr and Siegler observed in one of their subjects, who seemed to switch back and forth between weight and distance cues whenever the use of one led to an incorrect prediction on a conflict problem. It is noteworthy that this oscillation was not intentionally built into the model, but arises purely as a by-product of learning through discrimination and strengthening.

One detail of Klahr and Siegler's results which our transition model does not account for is their subjects' tendency to focus more on weights than on distances. In our model, at approximately the same time that DOWN-2 is constructed, a similar rule is created which includes a condition about greater distance instead of greater weight. While Klahr and Siegler found evidence for DOWN-2, they found none for the second rule. One can imagine introducing preferences into the model to focus attention on some attributes in favor of others, but unless one can explain where these preferences originated, they would provide no more explanation than labeling one dimension as more "salient" than another. Thus, our transition model does not account for every detail of Klahr and Siegler's results, but it does provide a plausible initial explanation of the transition between the observed stages.

CONCLUSION

In this paper we described a model of the developmental process on the balance scale task. Although parts of the model are necessarily limited to this domain, the learning mechanisms were implemented in a general way, and we are convinced that they could be used to explain developmental trends on other tasks. One drawback of the model is the speed with which it learns. The trials shown in Figure 2 required only 32 CPU seconds on a PDP-10 computer, while children take years to move from Stage 1 to Stage 3 behavior. However, the model was presented only with the relevant features of weight and distance, and if we had instead included other irrelevant information, its learning rate would presumably be reduced to a more reasonable level. Clearly, we have taken only a small step towards understanding the complex processes responsible for cognitive development, but we hope to have clarified one approach to studying this area.

REFERENCES

- [1] Baylor, G. W., Gascon, J., Lemoyne, G., and Pother, N. An information processing model of some seriation tasks. *Canadian Psychologist* 14 (1973) 167-196.
- [2] Young, R. M. *Seriation by Children: An Artificial Intelligence Analysis of a Piagetian Task*. Birkhauser, Basel, 1976.
- [3] Piaget, J. and Inhelder, B. *The Psychology of the Child*. Basic Books, New York, 1969.
- [4] Klahr, D. and Siegler, R. The representation of children's knowledge. In *Advances in Child Development*, Academic Press, New York, 1978.
- [5] Brazdil, P. Experimental learning model. *Proceedings of the Third AISB/GI Conference*, 1978, 46-50.
- [6] Anderson, J. R., and Langley, P. J. A learning system and its psychological implications. *Proceedings of the Sixth International Joint Conference on Artificial Intelligence*, 1979, 16-21.
- [7] Langley, P. Language acquisition through error recovery. *Cognition and Brain Theory* 5 (1982), 211-255.