Running head: MODELING DEVELOPMENTAL TRANSITIONS

Modeling Developmental Transitions on the Balance Scale Task

Hedderik van Rijn, Maarten van Someren, Han van der Maas

Department of Psychology,

University of Amsterdam,

Roetersstraat 15,

NL-1018 WB Amsterdam, The Netherlands.

rijn@swi.psy.uva.nl

Version of 2001/06/19, submitted to Cognitive Science

Please contact the first author before quoting this manuscript

Abstract

Cognitive development on certain tasks (e.g., the balance scale task) is characterized by periods of relative stable, rule-like but suboptimal behavior that occur in a fixed order and alternate with short transition periods. Many models have been developed to capture the developmental phenomena associated with the balance scale task. However, most of these models do not account for important phenomena as discontinuous transitions or rely on questionable assumptions, and none of the models is able to predict improvement in behavior <u>without</u> feedback. We propose a computational model that is implemented in ACT-R and is based on the evaluation of success of applied knowledge structures combined with a mechanism to construct new knowledge by searching for differences in the presented balance scale problems. This model accounts for the empirical phenomena, including learning without feedback as is common in developmental tasks.

Modeling Developmental Transitions on the Balance Scale Task

In many domains, cognitive development is conceptualized as a progression through a series of increasingly complex and accurate task specific strategies. Behavior on proportional reasoning tasks like the balance scale task is shown to adhere to this type of development: Longer periods of relatively stable performance alternate with short periods of instable performance in which behavior relatively sudden improves to a new phase. In a balance scale setting, a child is asked to predict the movement of a balance scale. Pegs are situated at equal distances from each other and the fulcrum. A number of equally heavy weights can be placed on one of the pegs at each side of the fulcrum. The balance will either tip to one side or remain in balance, depending on the number of weights and on the pegs at which these are placed. Typically, no feedback about the correctness of the answers is given during the experiment. This task is used to observe the strategies that children employ, to study the effect of training on the use of the strategies, and to study the development of one strategy into another (e.g., see Siegler, 1976, 1981; Siegler & Chen, 1998). Traditionally, six item types[1] are distinguished to categorize behavior: three simple item types and three conflict item types. The simple item types consist of: (1) Balance items having an equal number of weights, equidistant from the fulcrum; (2) Weight items with an unequal number of weights, equidistant from the fulcrum; and (3) Distance items with an equal number of weights at different distances from the fulcrum. For the conflict item types, the distance and weight dimensions indicate different predictions. Conflict items have unequal numbers of weights at different distances from the fulcrum: The side with the most weights has the smaller distance value, and vice versa. For (4) Conflict Weight items, the balance tilts to the side with the larger

[1]These six item types constitute 68% of all possible items. The remaining items have more weights at the side with the larger distance, and are therefore not suitable to categorize behavior. We will refer to these items as Weight/Distance items.

number of weights, for (5) Conflict Distance items, the tilt is to the side with the greater distance. For (6) Conflict Balance items, the effects of the two dimensions level out: the scale remains in balance. Studying the behavior of children on this task and, especially, comparing their performance on the different item types have provided a number of well-known and relatively stable empirical phenomena. These phenomena will be discussed in the next section. However, identifying phenomena does not provide an account of the underlying mechanisms of behavior. As argued by many authors (e.g., Boden, 1979; Anderson, 1987; Siegler, 1989), one needs to understand the mechanisms that cause the changes in behavior to obtain a causal explanation of development. In research that focused on explaining these mechanisms in the balance scale domain, a number of computational models have been proposed. However, none of the proposed models for balance scale behavior is able to explain the empirical phenomena to their full extent without depending on less plausible assumptions. In this paper, we will present computationally implemented explanations that account for the major empirical phenomena, and that do not depend on implausible assumptions.

Empirical Phenomena and Criteria

In this section, we will discuss the major empirical balance scale phenomena. This includes information that sheds light on the processes underlying changes in behavior. We will derive four empirical criteria based on these phenomena that computational models of balance scale behavior should meet.

Empirical Phenomena

Qualitatively Classifiable Behavior The behavior of children on the balance scale task is classifiable into qualitatively different "rules". The classification of behavior into different rules has been one of the cornerstones of the balance scale research. The assessment of rules is guided by the notion of homogeneous behavior on items from a particular item type. That is, the answers to a particular type of items are either all correct, or in all or most cases incorrect. This consistency criterion (Reese, 1989) is often cited as evidence for rule use. Within the balance scale domain, the consistency criterion is operationalized by Siegler's (1981) rule assessment methodology. Jansen and Van der Maas (1997) improved the rule assessment methodology by applying latent class analysis. This statistical method provides a more rigorous test of this criterion.

To improve performance, the existing rules have to be enhanced or replaced by better rules which implies some sort of transition between rules. During these transitions, behavior is less homogeneous as some items of a particular item type might already be answered correctly, whereas others are still answered according to the earlier, incorrect rule. How new rules are learned is not necessarily the same for all transitions. For example, most people need explicit instruction to be able to perform according to the most sophisticated rule, whereas transitions to earlier rules seem to come naturally.

The Rules Siegler (1976, 1981) identified four different patterns of behavior that can be characterized by Rules[2] (see Figure 1) that characterize children's behavior during different phases. Rule I: consider the amount of weights on each side: if equal, predict that the balance scale will remain level, if unequal, predict that the balance scale will tip to the side with greater weight. Rule II: also consistently predict that the side with the greater weights will tip, but if the weights are equal, the distance dimension is taken into account. Rule III: consider both the weight and distance dimension, but if one side has greater weight and the other greater distance, "muddle trough" or guess. Rule IV: consider both dimensions and compute the torques on each side by multiplication of the weights and distances of each side if the dimensions conflict.

[2]For clarity reasons, the Rules empirically identified by Siegler and others are written with a capital "R", whereas the ACT-R production rules, to be introduced later, are written with a small "r".

The unsystematicity of the behavior before Rule I makes a classification of behavior infeasible. The first three Rules are comparable because they consist of a number simple steps containing only comparisons. Rule IV contains an additional step that involves multiplication, which contrasts it with the simpler earlier Rules.

The six different item types have been carefully constructed to enable a categorization of behavior according to the Rules by analyzing the proportion correct answers per item type. For example, if most of the Balance, Weight and Conflict Weight items are answered correctly and most of the other items incorrectly, behavior is classified as the result of the use of Rule I. The period in which a child uses a particular Rule, is also referred to as a phase. For example, if behavior is according to Rule I, the child is categorized as in Phase I.

After Siegler's initial study, these Rules have been identified in numerous other studies. Besides confirmation of the Siegler Rules, a couple of new Rules were identified. These new Rules are mostly variants of Rule III, trying improve the classification of the "muddle through" behavior. A well-known example of a Rule III variant is the Addition Rule (proposed by Normandeau, Larivee, Roulin, & Longeot, 1989, as well as by Ferretti, Butterfield, Cah, & Kerkman, 1985). Children using the Addition Rule consider both weight and distance dimensions. If one side has greater weight and the other greater distance, they compare the sums of the number of weight and the distance.

Transition Phenomena To be able to describe the development from one Rule to another, detailed information is needed about the transition process. Information on these transitions has only recently become available. Jansen and Van der Maas (in press) describe the transition from Rule I to Rule II in terms of the cusp model from catastrophe theory (Thom, 1975). They found empirical evidence for a discontinuous phase transition by manipulating the distance dimension. The selection of distance was based on Siegler's (Siegler & Chen, 1998; Siegler, 1976, exp.3) research, which suggests that the availability of the distance dimension to the

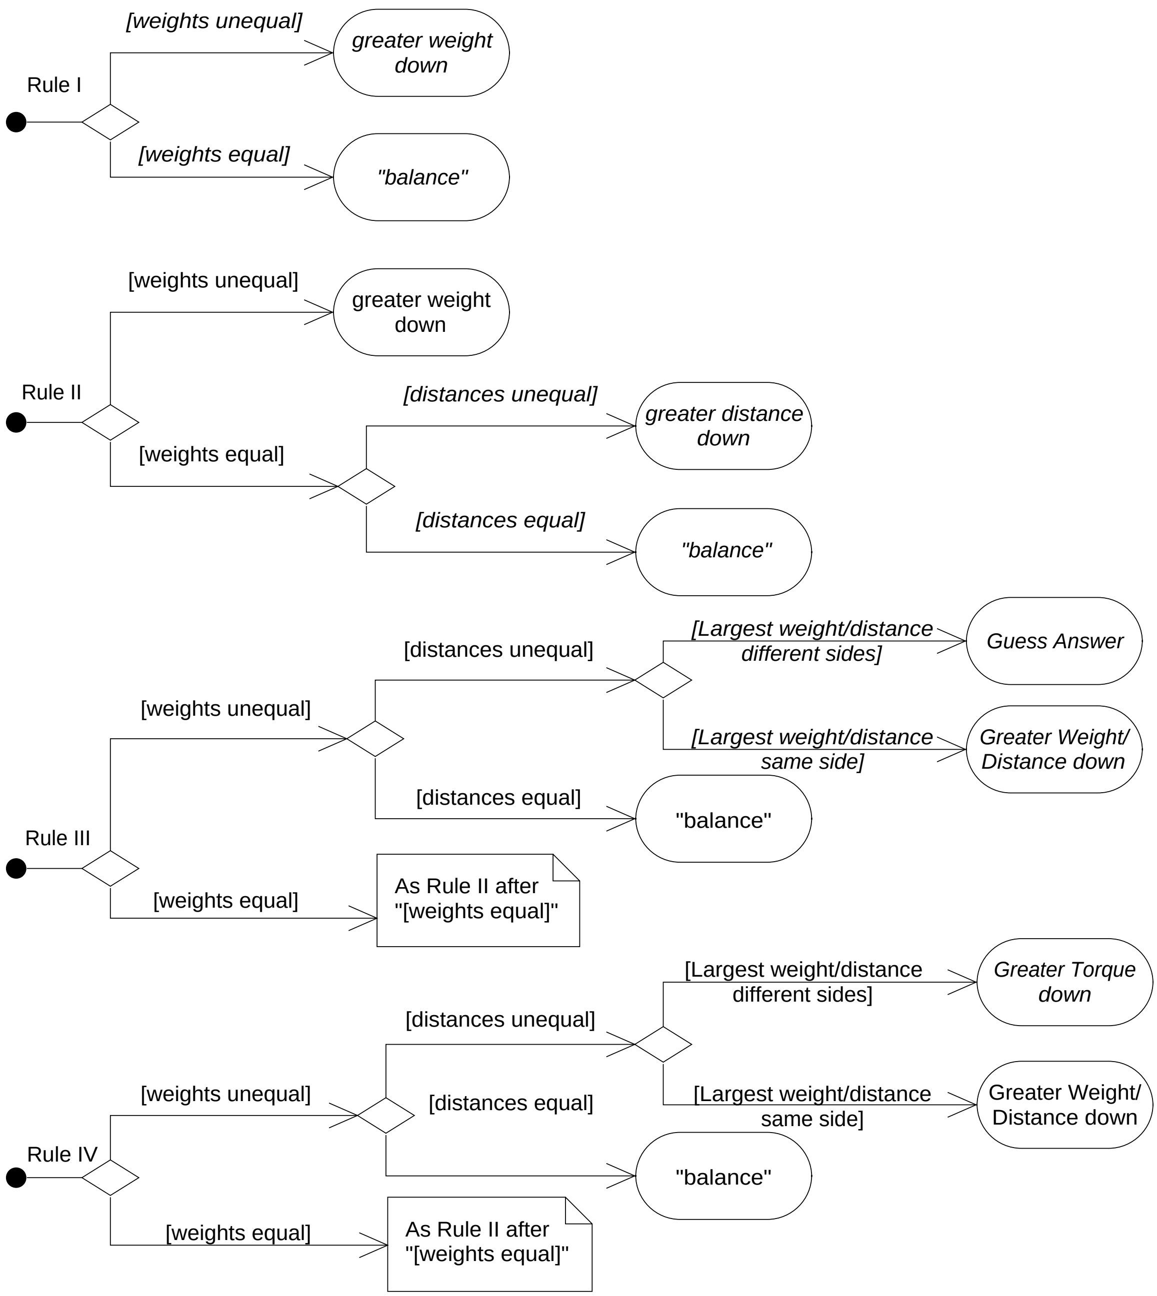

Figure 1. Rules that categorize behavior as identified by Siegler (1981).

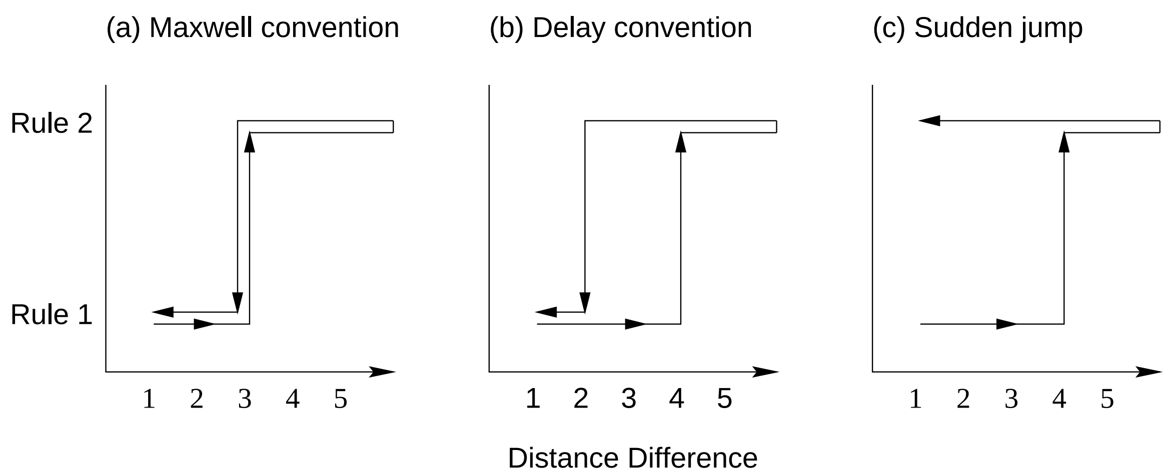

Figure 2. Three hysteresis patterns, x-axis denotes the difference in distances left and right of the fulcrum, the y-axis denotes the used Rule. Note: Children are presented a sequence of items which increasing distances first.

decision process mediates performance on the balance scale task. Jansen and Van der Maas presented children with sequences of distance items with the same weight configuration but different distance configurations. During the first part of the sequence, the difference between the left and right distance from the weights to the fulcrum increased with one per presentation; during the second part the distance difference decreased. As in most balance scale studies, the children were not given any feedback. During the sequence of increasing difference, a proportion of children were expected to change from Rule I to Rule II because of the increase in saliency of the distance and therefore the increased likelyhood of the distance dimenstion being available. The decrease was expected to cause a regress to Rule I in a proportion of the children. Their results speak in favor of a distance difference mediated phase transition from Rule I to Rule II.

Comparing the levels of the distance difference for the jump upward and the regress resulted in three different transition patterns. First, a group of children showed a jump upwards at the same level of the manipulated distance difference as the regress to the old level of behavior, which is called the Maxwell convention (Figure 2a, left panel). Second, children showing the delay convention (Figure 2b, middle panel) regressed to the old level of behavior at a smaller distance difference than the jump to the new level. Third, children showing a sudden jump (Figure 2c, right panel) did not regress to the old behavior, but remained at the new level of behavior for all levels of the distance difference after the jump to the new behavior.

Jansen and Van der Maas (in press) only studied the transition from Rule I to Rule II. Detailed empirical information on other transitions is not available.

Torque Difference Effect The homogeneity of answers for items of the same item type is questioned in research of Ferretti and Butterfield (Ferretti et al., 1985; Ferretti & Butterfield, 1992). They found evidence for a torque difference effect (TDE), that is, items with a large torque difference are more likely to be answered according to a more advanced Rule than items with a smaller torque difference. This implies non-homogeneity of behavior within item types. However, reanalysis of their data (Jansen & Van der Maas, 1997) shows that the torque difference effect only occurs for items with extreme torque differences. Therfore, it is concluded that the behavior is homogeneous for non-extreme torque difference levels and that a torque difference effect occurs for large torque differences.

Empirical Criteria

Based on these phenomena, the following four empirical criteria will be used to judge to what extent the computational models explain the observed data.

EC1. Rule-like Behavior The behavior of the model should be classifiable into Rules. Ideally, this classification can be conducted by latent class analysis or by inspecting the structure of the models[3]. As rule-like behavior implies transitions, the model must explain transitions

[3]Some models directly implement the Rules as specified in Figure 1. Classification by latent class analysis is superfluous for these models.

from one Rule to another.

EC2. Rule sets A complete model of balance scale behavior should at least include the four Rules as identified by Siegler (1976) and the Addition Rule. Moreover, Rule IV should not be learned by an approximation of multiplication (e.g., an extensive set of item-answer mappings), but by calculation of the torque. The Rules should become available in a fixed order, independent of items presented to the model.

EC3. No Feedback and Transition Patterns First, Jansen and Van der Maas (in press) showed that transitions are possible on series of items with increasing saliency of one of the dimensions, without feedback about the correctness of answers. Second, if the saliency of the distance dimension is manipulated, the model should show the three transition patterns as presented in Figure 2.

EC4. Torque Difference Effect for Large Torque Differences A model should explain the Torque Difference Effect for large torque differences, but should not show a torque difference effect for small difference values.

Computational Models

The underlying principles of the computational models of balance scale behavior range from a static set of simple production rules to complex, self-expanding clusters of information processing units. Because we do not want to enter the discussion on underlying principles of computational models, we will primarily discuss the models in terms of their adherence to the empirical criteria (EC1 to EC4).

However, we will test the models against one computational criterion: each model's domain related assumptions should be plausible and inspectable. This holds especially for the assumptions related to the representation of the balance scale items, the prior knowledge available to the model in different stages of behavior, and the results of the proposed learning mechanism. Moreover, the model must not be too sensitive to minor modifications in the assumptions. Because of the emphasis on the empirical phenomena, we will only discuss this criterion if the model adheres to the empirical criteria.

First, we will describe the symbolic production rule models of balance scale behavior. Second, we will discuss the neural network models. Third, we will discuss a model based on a decision tree mechanism.

Production Rule Models

We discuss two production rule models[4]. The first computational model of balance scale behavior is Klahr and Siegler's (1978) production rule model. They showed that different sets of static production rules were able to capture the observed Rules. The model did not specify a transition mechanism and it therefore fails the first empirical criterion (EC1) by not modeling development.

Sage and Langley (Sage & Langley, 1983; Langley, 1987) added a transition mechanism. Their model starts with a set of production rules that produce random answers. The model learns by discrimination: if the model predicts an answer that is found to be incorrect, a new production rule is created based on an analysis of the differences between the incorrectly answered item and the last item for which the production rule performed correctly. This new production rule has an additional condition that discriminates between the two items and the appropriate answer. Performance is kept stable for some time by parameters that determine when a production rule is replaced by another rule. This model shows rule-like behavior and transitions (EC1). However, the model does not learn Rule IV because it is not equipped with the necessary knowledge of multiplication. Another item is that the discrimination method often leads to production rules that

[4] A often cited production rule model of balance scale behavior is Newell's (1990) SOAR model. However, the sparse description of this model does not provide enough information to evaluate this model.

are not found in human behavior. The model therefore fails EC2. Only if knowledge about the correctness of a prediction is available, the model learns new Rules and therefore fails the no-feedback criterion (EC3). Since the model only checks whether weights and distances differ or not, it does not show the Torque Difference Effect for extreme torque differences (EC4).

Connectionist Models

McClelland (McClelland, 1989, 1995) presented a connectionist model of balance scale behavior that learns by back-propagation of feedback. The model consists of two output nodes representing the predicted answer, and five input nodes for weights and for distances per side of the scale. The input and output nodes are connected by separate hidden layers for the distance and weight dimensions. To explain the use of the weight dimension before the distance dimension, McClelland assumes that weight items are (initially) more frequent than distance items.

The first empirical criterion poses a problem for this model. McClelland tested the model using Siegler’s rule assessment methodology, claiming a successful fit to human data. However, reanalysis of the model’s behavior using latent class analysis has shown that the behavior of the model cannot be described by a set of distinct rules unlike human data (Jansen & Van der Maas, 1997). Therefore, the model fails EC1. With respect to EC2, the model is not able to learn stable Rule IV behavior. It sometimes regresses after a period of stable behavior in a more advanced phase. Because the model is only able to learn by feedback-dependent back-propagation, it fails the no-feedback criterion (EC3). With respect to EC4, the model does show the Torque Difference Effect (McClelland, 1995). However, Jansen and Van der Maas (1997) argued that this model fails to show homogeneity within item types at all. Furthermore, Raijmakers, Van Koten, and Molenaar (1996) showed by reanalysing the model that it fails to show qualitative transitions between the Rules.

Shultz and his co-workers (Shultz & Schmidt, 1991; Shultz, Mareschal, & Schmidt, 1994; Shultz, Schmidt, Buckingham, & Mareschal, 1995) modeled balance scale behavior with cascade-correlation networks. The central tenet of cascade-correlation is that if the back-propagation-like learning mechanism does not reduce the error between predictions and feedback fast enough, new hidden nodes are incorporated into the model. Here we will discuss the model as presented in Shultz et al. (1995). This model has five input units for each side and two output units. The distance dimension determines which of the input units it activated, and the weight dimension determines the amount of activation. The model is pre-trained with weight items to model the initial preference for weight.

As in McClelland's model, the behavior of the cascade correlation network is categorizable into Rules using Siegler's rule assessment methodology. Regrettably, no latent class analysis has been applied to the model's data. However, because the learning mechanism of the cascade correlation network is comparable to that of McClelland's model, it is not likely that the overall pattern of gradual adaptation is different. Our impression of gradual adaptation is reinforced by Shultz' (1994, p. 81) remark that: "The cascade correlation model suggests that balance scale [..] transitions are soft and tentative rather than abrupt and definitive". It is therefore likely that the results of a latent class analysis are comparable to the analysis of McClelland's model (i.e., fail EC1). This model is able to learn all four Rules in the correct order (EC2). However, as no torque is being calculated to come up with an answer, the model only partly (i.e., only for the trained domain) adheres to Rule IV behavior. With respect to the no-feedback (EC3) and the Torque Difference Effect/homogeneity (EC4) criteria, the same comments hold as for McClelland's model.

Decision Tree Model

Schmidt and Ling (1996) presented a model based on decision tree learning. This model learns behavior on the balance scale by constructing decision trees. This mechanism involves explicit extraction of rule-like regularities from a set of data. At the root of the tree is the most predictive condition, one level below are the next most predictive decisions, etcetera. For example, a decision tree for Rule II (Schmidt & Ling, 1996, p217) starts with "Are weights and distances equal?” If both are equal, the model answers “balance”, otherwise, the next step is “Which side has the most weights?”. During the construction of the tree, new conditions are added until the tree reaches a pre-specified “error tolerance” threshold. Note that although alternative decision tree algorithms exist, the algorithm used by Schmidt and Ling requires that a complete training-set is available to the model when constructing the tree. It is not able to learn incrementally but Schmidt and Ling claim that comparable trees are constructed when incremental learning algorithms would have been used. Therefore, the different states of the model are better described as “different snapshots of development”, instead of illustrative for “continuous development”. The input to the model deviates from the input to the other models. Not only the numbers of weights and distances are presented to the model, but also (1) a binary attribute representing whether both weights and distances are equal, (2) the difference in weights between both sides and (3) the difference in distances between left and right.

The model adheres to the first empirical criterion, as the Rules can be identified by inspecting the constructed decision trees and the incremental version of decision tree mechanism is probably able to simulate the transitions (EC1). The model also learns the Rules in the right order (EC2). However, because no knowledge of multiplication is represented in the model, it does not calculate the torque which is necessary for correct Rule IV behavior. Instead, the model’s approximation of Rule IV is based on extensive item-to-answer mappings. EC3 is not easily evaluated as the model does not learn new Rules incrementally. Therefore, the transitions from one Rule to another are by definition sudden and abrupt. Because the construction of a new tree is only based on detecting regularities in the mapping of items to the associated answers, improving behavior without feedback is impossible. Hence, the model does not fulfill the no-feedback criterion (EC3). The model is able to construct different Rules for different difference intervals because the model has access to the differences in weights and distances. Therefore, it is able to model the Torque Difference Effect. Even more, because the distance difference has to exceed a certain value before the Torque Difference Effect occurs, the homogeneity of items with a smaller torque differences is preserved (EC4).

Based on the empirical criteria, the decision tree model simulates development on the balance scale task relatively well. However, one of the aspects of the computational criterion is the input representation. To be able to meet criteria EC1, EC2 and EC4, the differences between left and right values for weight and distance had to be included. This assumes that this information is always and automatically available to a child solving a balance scale problem. To be able to perform at Rule IV level, the decision tree model has to explicitly include all possible balance scale items that are not solvable by comparisons. This results in a complex and extensive decision tree. It might be questioned whether these input representation and the decision tree structure assumptions are plausible.

Conclusions

As discussed above, although all models provide an interesting point of view on the process of balance scale development, none is able to model the empirical data to full extent. The main problems are that learning without feedback are not captured in the presented models. Furthermore, the non-symbolic models have difficulty simulating rule-like behavior whereas the symbolic models require less plausible assumptions or are not able to correctly model the development of new Rules. Below, we present two models of balance scale behavior using the ACT-R cognitive architecture which overcome most of these problems.

Models of Balance Scale Behavior

In this part of this paper, we will present two models of cognitive development on the balance scale task. The first model uses feedback on its answers to pass through the sequence of behavior. The second model simulates the construction of Rules in the absence of feedback. As we will show, together these two model meet the empirical critera. Both models are implemented in ACT-R (Anderson, 1990; Anderson & Lebiere, 1998 and see Lebiere & Anderson, 1998 for an earlier demonstration of the application of ACT-R to cognitive development). ACT-R stores knowledge in two types. Factual knowledge is stored in the form of declarative chunks, whereas actions are stored as production rules, which take the form of IF-THEN rules. Each of these types of knowledge has an associated value that expresses the empirical usefulness of that piece of knowledge. This value is related to the key assumption of ACT-R: the process of rational analysis. An ACT-R model constantly re-evaluates the usefulness of its knowledge by rational analysis. If knowledge proves to be relatively successful, it will be selected for usage more often, whereas less successful knowledge will be less often selected. Ultimately, this decrease renders the knowledge impossible to access if the associated value drops below a threshold. Before discussing the subsymbolic rational analysis, we will discuss the chunks and production rules at the symbolic level.

Each chunk belongs to a particular type. The type determines the structure of that chunk: the slots that can be filled with variable information. Production rules can request the retrieval of a particular type of chunk. Chunks play an important role in the guidance of an ACT-R model. At start, a chunk is tagged as goal. This goal chunk determines which production rules can be applied.

Each production rule specifies to which goal type (i.e., chunk type) it belongs. The production rule can only be applied if its type is equal to the type of the goal chunk. Generally, the IF-part consists of constraints on the current goal (e.g., a production rule to say the answer to a problem can only be applied if the answer is already available in the current goal) and a retrieval of a chunk (e.g., a phonetic pattern associated with the answer). The THEN-part of a production rule expresses the actions to be taken if the constraints are satisfied and the requested chunk can be retrieved. Often, a complex task can be divided into subtasks that have their own set of applicable production rules. Within a model, this is achieved by sub-goaling (e.g., the task to answer a simple arithmetic problem can be divided into a subgoal to find the answer to a problem and a different subgoal to say the found answer). This process is initiated in ACT-R models by pushing a new goal on the goal stack which then becomes the current goal.

The subsymbolic value of chunks associated with rational analysis is their activation. Each chunk has an activation which is determined by the number and recency of past references to that chunk. The activation of a chunk is determined by the following formula: B = $\ln(\sum_{j=1}^{n} t_j^{-d})$. The summation is over all previous references to this chunk, with t_j being the time in seconds between reference j and present and d being the decay rate (often fixed at .5). If a model needs to retrieve a chunk, and if more than one chunk reaches threshold and satisfies the constraints of that retrieval request, the chunk with the highest activation will be retrieved.

For production rules, the value used for rational analysis is the expected gain. The expected gain of production rules is determined by a combination of the proportion of successful completions of goals in which that production rule played part and the total costs associated with that production rule. This is expressed in the following formula: Expected Gain = EG = PG - C. The P stands for the proportion of successful completions, G is a constant value representing the value of the goal, and C stands for the costs associated with the production rule. As with chunks, if more than one production rule meets the current constraints, the production rule with the highest expected gain is selected for execution. With regard to subgoaling, the expected gain variables of a production rule depend on that production rule's success and on the subgoals it initiates. For example, if the retrieved answer to an arithmetic problem is incorrect, the expected gains of the production rules that were used to express the answer are not affected because these were used within the context of a subgoal.

If a production rule explicitly removes the goal chunk from the stack of goals, the goal is completed successfully. However, if at a certain point no production rule above threshold matches the goal, the architecture removes this goal from the goal stack and all applied production rules are ascribed a failure which lowers the expected gain.

A recent addition[5] to ACT-R (Taatgen & Anderson, submitted; Taatgen, 2000) proposes a

[5] This addition is not yet part of the current ACT-R version, but is incorporated in the forthcoming ACT-R 5.0. Although the current version of ACT-R is also able to construct new production rules, the new mechanism constructs new rules using a more elegant and principled method.

proceduralization mechanism that constructs new production rules. Instead of supplying the model with production rules (implying that these rules have been constructed outside the scope of the model), the proceduralization mechanism uses declarative representations of actions. Production rules are learned in one of two ways: (1) declarative descriptions of actions are transformed into production rules, (2) existing rules are merged into a new, specialized production rule.

With respect to the first method, during the process of reasoning with the declarative representations new production rules can be constructed on a probabilistic basis. This reflects the notion that automatizing behavior is a side effect of performing (an earlier form of) the same behavior. So, given enough experience with a declarative representation, it is likely that a similar production rule is created. Because the cost (and therefore the expected gain) of one production rule is lower than the cost of reasoning with the declarative representation and performing the same actions, the new production rule will take over. The second method constructs a new production rule when two rules are executed after another with the first rule retrieving a chunk. A new, specialized production rule is formed that combines the two, eliminating the step that retrieves of the chunk. That is, instead of retrieving a chunk, that chunk's information is hard coded into the production rule. For an extended overview of proceduralization, refer to Taatgen and Anderson (submitted).

A Feedback ACT-R Model of Balance Scale Behavior

The model that we present in this section consists of a set of production rules that resemble the original Rules. We show that the standard ACT-R expected gain learning mechanism successfully produces the correct sequence of transitions.

The Rules presented by Siegler are presented in the form of decision trees (see Figure 1). For example, Rule I is represented as: if the weights are equal, answer "balance", otherwise give as answer the side with the most weights. To describe this Rule in terms of IF-THEN rules, two IF-THEN rules have to be created: (1) IF the weights are equal THEN say balance, (2) IF the weights are unequal THEN push a subgoal to give as answer the side with most weights.

Table 1 shows all production rules derived from the decision trees[6]. As not all production rules are used for all Rules, the first column specifies in which Rule the production rule occurs. The first three production rules in this table (1-3) are relatively simple production rules that reflect the initial steps of starting with a problem and retrieving the number of weights on the left and the right side. Then, either production rule 4a or 4b is selected for execution, depending on the equality of the retrieved weight values. For the transition to Rule II, production rule 4a has to be replaced with production rule 4a-1 which starts a new sequence of production rules that process distance information (4a-1;4a-2;4a-3a;4a-3b). The next transition is signaled by the replacement of production rule 4b by 4b-1. This again leads to a new sequence of production rules which process distance information (4b-1;4b-2;4b-3b;4b-3c). The last transition is the replacement of production rule 4b-3c with the force rule 4b-3d.

To explain the order of Rules in development, the model has to explicate how some production rules are used in earlier behavior (e.g., 4a and 4b), whereas other production rules are not used until later in development (e.g., 4b-1 and 4b-3d). As introduced before, all production rules in ACT-R are subjected to expected gain evaluation. The production rule which has the highest expected gain, is selected for execution. This value is calculated by EG = PG - C. As G is kept constant (at the default setting of 20), P and C are the values that can be used to explain development.

The costs of production rules ("C") are assumed to be determined externally because of application of the rules in other tasks or domains. It is unlikely that tasks like comparison, addition or multiplication are first used within the context of the balance scale task. Obviously, the costs of

[6]With respect to Rule III, this overview deviates from the original decision tree specifications. As empirical studies (Jansen & Van der Maas, 1997) demonstrated the use of the Addition Rule instead of the original Rule III, the model presumes the use of the Addition Rule as specified by Normandeau et al. (1989).

Table 1

Production rule representation of balance scale Rules.

Rule	Number	Production rule description
I-IV	1.	Sets a goal to search for an answer based on properties of the presented problem.
I-IV	2.	Selects the weight property as initial property to base an answer on.
I-IV	3.	A subgoal that retrieves the values of the previously selected property from the balance scale.
I	4a.	IF the weights are equal, THEN give “balance” as answer.
II-IV	4a-1.	IF the weights are equal, THEN selects the distance property for further processing.
II-IV	4a-2.	See production rule 3.
II-IV	4a-3a.	IF the weights and the distances are equal, THEN give “balance” as answer.
II-IV	4a-3b.	IF the weights are equal and the distance values are retrieved, THEN base the answer on the distances.
I-II	4b.	IF weights are not equal, THEN select as answer the side with most weights.
III-IV	4b-1.	IF the weights are not equal, THEN select the distance property for further processing.
III-IV	4b-2.	See production rule 3.
III-IV	4b-3a.	IF the distances are equal, THEN base the answer on the weights.
III-IV	4b-3b.	IF the weights and distances have a larger value on the same side, THEN give as answer that side.
III	4b-3c.	IF weights and distances are retrieved, THEN base the answer on the addition of weights and distances.
IV	4b-3d.	IF weights and distances are retrieved, THEN base the answer on the multiplication of weights and distances.

these operations decrease with age. However, the costs of all these operations decrease. Because only the differences between costs are relevant, We make the simplifying assumption that the C value is static during development. Although the cost of production rules is set in absolute values (by default expressing time in seconds), within this model the costs must be regarded relative to the other costs. We used the following relation to estimate the costs of the production rules: cost of comparison resulting in an equality < cost of comparison resulting in an inequality < cost of addition < cost of multiplication. (The relations used are 0.5:1:3:4, although the model is robust to minor modifications of these values.)

The other parameter ("P") of the expected gain formula reflects the prior success-rate of the associated production rule. Although particular methods (e.g., multiplication) might be used before in other domains, their success-rate is unknown for the new domain of balance scale behavior. As this holds for all production rules, we set the P value of all production rules to the same value: P = .875 = 175 successes / (175 successes + 25 failures). This "initial experience" reflects the confidence in the yet unused rules, but claiming that the untested production rule will predict perfectly[7]. Without setting these "experiences", the first balance scale item could cause a sudden shift in the P value. The qualitative behavior of the model is insensitive to modifications in these numbers, as long as the initial number of successes is smaller than the number of failures.

The P value is the main source of development in this model. If the model is presented with randomly chosen balance scale items, the P values of the production rules will be adjusted to reflect the empirical evidence. The number and percentage of correct responses per Rule is presented in Table 2.

As can be seen in this Table, the earlier production rules have a lower success-rate than the later production rules. Because of this difference in success-rate, production rule 4a-1 (20% correct) will be replaced by production rule 4a-1 (100% correct). In a similar fashion, 4b (78%) will be replaced by 4b-1 which lead to an answer given by either 4b-3a, 4b-3b or 4b-3c (91% correct). Finally, the not 100% successful production rule of Rule III (4b-3c) will be replaced by

[7] See, for example, Lovett and Anderson (1996) for an illustration and discussion of using the P parameter to model behavior.

Table 2

Number and percentage of correct responses per combination of Rule and balance scale type.

	Balance Scale Items							Total	
	B	D	W	CB	CD	CW	WD	Correct Responses	
Rule I	25	0	100	0	0	88	200	413	66%
weight equal, 4a	25	0	∅	∅	∅	∅	∅	25[b]	20%
weights unequal, 4b	∅	∅	100	0	0	88	200	388[c]	78%
Rule II	25	100	100	0	0	88	200	513	82%
weight equal, 4a-1	25	100	∅	∅	∅	∅	∅	125[b]	100%
weights unequal, 4b	∅	∅	100	0	0	88	200	388[c]	78%
Rule III									
Addition,[a] 4b-3[a-c]	25	100	100	20	64	64	200	573	91%
Rule IV, 4b-3d	25	100	100	24	88	88	200	625	100%

Note. The numbers after Rule IV and the subsets of the other Rules refer to the production rules in Table 1 that are directly related to this behavior. The absolute numbers assume a five $\times$ five balance scale configuration. Different configurations lead to both different absolute and relative values. B=balance item, D=distance item, W=weight item, CB=conflict balance item, CD=conflict distance item, CW=conflict weight item, WD=weight/distance item.

[a] As described in Normandeau et al. (1989).

[b] Of a total of 125 items with equal weights.

[c] Of a total of 500 items with unequal weights.

4b-3d (100%).

<u>Simulation</u>

To test the behavior of the model, randomly selected balance scale items were presented to the model. This was done by setting a goal to solve a balance scale problem, and adding the appropriate weight and distance chunks to declarative memory. The training set was not biased to weight items nor to balance items (as sets in some of the previous studies were). Each time an answer to a balance scale problem is given, the expected gains of the involved production rules are updated. Development takes place because of this updating. The expected gains of some main production rules can be used to illustrate the current state of the model. Because production rules like the first three production rules in Table 1 are used during all phases, they do not affect the behavior of the model. Another pair of production rules of which the expected gain is less informative, are 4a-3a, 4a-3b, 4b-3a and 4b-3b. These production rules have to be triggered by one of the two production rules (4a-1 and 4b-1) that shift the attention to the distance feature. As the answers given by these production rules are all correct, these production rules will not be replaced by other rules. Moreover, their expected gain will increase to an asymptote of P=1. Therefore, only the following production rules are informative about the current state of the model: 4a, 4b, 4a-1, 4b-1, 4b-3c and 4b-3d. As soon as 4a is replaced by 4a-1, the model behaves according to Rule II. When 4b is replaced by 4b-1, the model is able to include both weight and distance in its decision process, leading to either the use of 4b-3c (Rule III, Addition) or 4b-3d (Rule IV, Multiplication).

The expected gains of these production rules are plotted in Figure . For clarity, the x-axis is plotted in logarithmic scale. Initially (upto about 100 training items), the expected gains of only two production rules are modified. The topmost expected gain (solid line) is associated with production rule 4a (“weights equal, answer balance”), whereas the lower line (dashed line) is the expected gain of 4b (“weights unequal, answer heaviest side”). Although the first rule starts out with a higher expected gain because it contains a comparison resulting in an equality, the probability of success of that rule is only .20. Therefore, the expected gain of this production rule decreases fast, which is visible in the steeper negative slope of the plotted expected gain. This fast

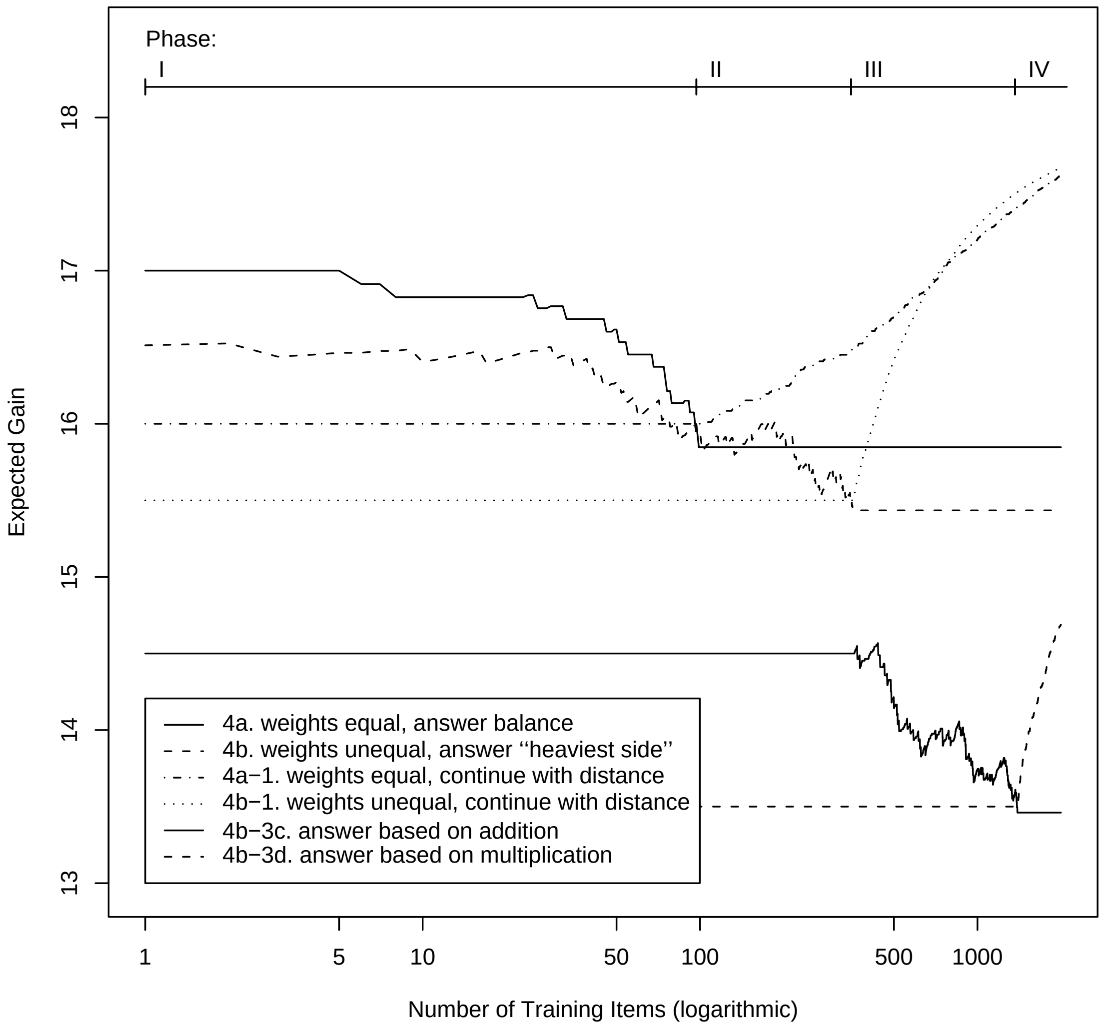

Figure 3. Expected gains of selected production rules. See text for a discussion of these rules.

decrease causes the replacement of this production rule by 4a-1 (dot-dashed line, “weights equal, continue with distance”) sooner than 4b is replaced. The sequence which is initiated by 4a-1 performs flawlessly, resulting in a slow but steady increase of the P parameter to an asymptote of P=1. However, this increase is relatively slow because this production rule is only used in approximately 24% of the cases.

If a new production rules causes an error just after it was first used, a period of instable performance occurs in which the new and old production rule are used in turns. Because the new “continue-with-distance” rule yields perfect answers, no instable performance is shown during the transition to Rule II. After all, if the weights are equal, the distances determine the tilt of the balance scale.

As the expected gain of the weight-unequal production rule 4b (dashed line) decreases, it drops below the expected gain of 4b-1 (dotted line, “weights not equal, continue with distance”). This triggers the model to take both weights and distances into account.

After this production rule has been executed, the model has access to both weight and distance information. The majority of items can be solved by the “cheap” (considering the costs, “C”) comparison-based production rules (4b-3a and 4b-3b). However, if these do not result in a correct response, the model tries to combine the information on weights and distances. The first Rule that implies a combination is the Addition Rule (production rule 4b-3c). Because the Addition Rule performs relatively poor for the items it applies to (148 out of 200 correct or 74%), it will be replaced by the force production rule, 4b-3d, which yields perfect answers.

Discussion of the Simulation Results With respect to the first empirical phenomenon, EC1, the above presented model does show rule-like behavior as can be seen in Figure . If performance has reached a certain level, it stays at that level for a while before improving to the next level.

Obviously, the model adheres to Siegler’s Rules as the decision trees were used to construct the production rules. This was confirmed by latent class analysis of the model’s responses. The main question was whether these Rules were learned at all and whether they occur in the correct sequence. Because all Rules are learned and learned in correct sequence, the model satisfies EC2. The model fails EC3 as the learning of new production rules depends on feedback. The model as presented above does not have any information about the size of the differences on the weight and distance dimension. Therefore, it is not able to show any torque difference effects and fails EC4.

Another issue to be raised is the computational criterion. The above presented model is equipped with all necessary production rules from start. Although the more complex production rules are not used until later in the simulated development, the model does not explain how these new production rules are instantiated.

Although the model fails both EC3 and EC4, this model does show that if one considers the Rules to be increasingly difficult, exposing the internal representation of the Rules to randomly selected items is sufficient to show transitions from Rule I to Rule IV. In ACT-R terms, the initial costs and the adaptation of production rules parameters explain the order in which development takes place. On a more general level, this model shows how gradual adaptation is able to exhibit performance that is characterized by periods of stable behavior, separated by sudden jumps.

Below, we present a model of the transition from Rule I to Rule II based on detailed empirical knowledge. Because this model focuses on a transition, the main criteria are EC3 and EC4 and the computational criterion of plausibility.

Modeling Transitions without Feedback

To explain what causes a transition to a later Rule, one has to explain how an earlier Rule can be replaced by a later Rule, and how new knowledge, necessary for this later Rule, is acquired. Above, we discussed how the evaluation of success of production rules explains the development through the identified phases. This approach presupposed that all necessary knowledge was readily available to the model. In this section we will discuss a model that is able to construct new production rules to solve balance scale items. To construct a model that explains transitions rather than merely shows transitions, detailed information is needed about the transitions. Recently, two studies (Siegler & Chen, 1998; Jansen & Van der Maas, in press) have been reported that studied phenomena related to the transition from Rule I to Rule II. Therefore, this model will focus on that particular transition.

As discussed earlier, children who are attributed to Rule I often do not notice the distance dimension. Siegler and Chen (1998) asked children to reconstruct a balance scale problem after they studied it for 10 seconds. In correspondence with earlier findings (Siegler, 1976), children attributed to Rule I were less likely to reconstruct the distance dimension correctly than Rule II children. A related finding is reported by Jansen and Van der Maas (in press). They showed that enhancing the saliency of the distance feature increased the proportion of items that were answered correctly. Possibly, the previous level of performance was at least partly caused by the inability to incorporate distance in the decision process. This suggests the following mechanism: Answers to balance scale problems are based on a search for differences in related properties.

This leads to the following hypothesis to explain the transitions: "A transition is only possible if a necessary property becomes available". Although not much is known about the performance before Rule I, the main difference is that Rule I children are able to systematically incorporate the weight dimension into their decision process. Therefore, this same mechanism might apply for the transition to Rule I.

Because no feedback was given in the studies reported by Jansen and Van der Maas (in press) and Siegler and Chen (1998), expected gain modifications like in the previously presented model cannot explain the transitions from Rule I to Rule II. Jansen and Van der Maas manipulated the saliency of the distance dimension by increasing the distance difference between the left and right side of the balance for consecutive items. Siegler and Chen presented subjects with the answers to a number of distance items and asked for an explanation of the observed behavior. The manipulations in both experiments drew attention to the relevant dimension, which increased the probability of incorporating that feature. This suggests that the dimensions themselves have to be encoded before the values of those dimensions can be incorporated into the decision process. For example, if a child is presented a balance scale problem, the first decision has to be what information of that balance scale will be used to base an answer on. It is possible that certain features are not noticed as being potentially useful because their internal representation (a chunk, in ACT-R terms) is not available to the child at that moment, for example because the chunk's activation is below threshold. We modeled the stressing of the distance dimension in abovementioned studies by an increase in the associated chunk's activation level. If this increase raises the activation above a threshold, the dimension can be used and incorporated into the decision process.

Besides explaining why performance is stuck or develops, a model also has to explain the construction of the knowledge on which new performance is based. Hereto, we will use the new ACT-R proceduralization mechanism. We will show that a model that combines the threshold-based retrieval of features and the proceduralization mechanism is able to explain development from Rule I to Rule II, even without feedback.

The Structure of the Model The model starts without <u>any</u> production rules about how to answer balance scale problems. Only general purpose production rules are available. These production rules are concerned with interpreting declarative representations (for example, selecting a next applicable action, carrying out retrievals, etc.). This knowledge is completed with a production rule that gives an answer to a presented problem (i.e., guesses) if no further steps can be taken. The following declarative representations of actions are available:

- If an answer has to be given to a difference problem, try to retrieve a yet unused feature to solve the problem
- After selecting a feature, try to retrieve the associated values
- If the retrieved values are equal, then start again
- If the retrieved values are unequal, base the answer on the observed difference

This declarative representation presupposes an important role for the concept of inequality. We assume that children start searching for an inequality if presented with balance scale items. Because they are instructed to predict which side goes down, it is likely that they will try to find which dimension differs between the two sides. Where other computational approaches presented the model a higher proportion of weight items, we assume a preference for the weight feature expressed in an additional activation for the chunk representing the weight dimension. Without this preference, the weight feature would be below the retrieval threshold like the distance feature. This is based on the assumption that children "know" that weight is important if they have to predict which side goes down.

The model is presented with randomly selected balance scale items. The above presented declarative representations are being interpreted to answer these problems. As a side effect of this interpretation, new production rules are formed. The following set of rules[8] emerges:

1. *Start-And-Retrieve-Weight-Values*: Starts the problem solving process, and directly retrieves the weight values.

2a. *Give-Answer-If-Weights-Are-Unequal*: Selects which answer will be given, based on the side with the larger number of weights.

2b. *Start-Search-For-Another-Feature-If-Values-Are-Equal*: If the retrieved values are equal, start over by searching for another feature.

This set of production rules is functionally equivalent to the production rules (sequence 1;2;3;4a/4b) associated with Rule I of the basic model presented earlier. One of the differences in these production rules, compared to the previous model's production rules, is the removed subgoal for retrieving the weights. The production rules of the new model are specialized by the proceduralization architecture to automatically retrieve the weights if a balance scale problem is

[8]Note that the names of these production rules are <u>not</u> constructed by the model but are assigned post-hoc for explanatory reasons.

encountered. If the weight values are equal, the model searches for another feature. However, we assume that, initially, the activation of the distance feature is too low to be retrieved. This renders step 2b of this model to be functionally equivalent to the weights-equal production rule (step 4a) of the previous model. Because no difference in features is found, the production rule that gives an answer if no further steps can be taken answers “balance”. As a result, the model’s behavior is comparable to the behavior of children in Rule I.

If the activation of the distance feature is increased, for example because of extended cueing of the distance dimension, the model starts to incorporate the distance feature in its problem solving process. In the same fashion as it learned its own production rules for Rule I, the model learns new production rules for Rule II behavior. Note that this improvement in behavior does not rely on the explicit addition of new knowledge. The same knowledge that leads to the transition to Rule I leads to the transition to Rule II. After the proceduralization mechanism has proceduralized the representations and merged and specialized the production rules, the following production rules result:

1. *Start-And-Retrieve-Weight-Values*: As above.

2a. *Give-Answer-If-Weights-Are-Unequal*: As above.

2b-1. *If-Weight-Equal-Continue-With-Distance*: If the weights are equal, note in the goal that the answer should be based on the distance feature.

2b-2. *Retrieve-Distance-Values*: Retrieve the distance values from the presented balance scale.

2b-3a. *Answer-If-Distances-Are-Unequal*: Base the answer on the values of the retrieved distance values.

2b-3b. *Start-Search-For-Another-Feature-If-Values-Are-Equal*: As above.

As with the model of Rule I behavior, this state of the current model is highly comparable to the basic model, with two modifications/extentions: (1) the retrieval of the values of the dimensions are specialized and (2) the search for another feature is functionally equivalent to a balance answer because no other features can be found. This shows that the assumption of a proceduralization process and simple prior knowledge is sufficient to simulate development from Rule I to Rule II.

Simulations: Transition Patterns To test the model, we compared the model's behavior to individual patterns acquired in the empirical study of Jansen and Van der Maas (in press). The model is presented with the same sequence of items as in this study. Like the children in the empirical experiment, the model receives no feedback about the correctness of its answers. All items have two weights on each side of a five pegged balance. In the first item, the weights were placed on the first peg left of the fulcrum and on the second peg right of the fulcrum, resulting in a distance difference of one. In the next four items, the weights on the right side of the fulcrum were shifted one peg to the right. When the distance difference reached the maximum of four, the sequence was reversed.

As discussed earlier (see Figure 2) Jansen and Van der Maas observed three types of transitions when the salieny of distance was manipulated: (1) A Maxwell convention pattern: when the jump to Rule II occurs at the same distance difference as the regress to Rule I; (2) a delay convention pattern: the jump to Rule II occurs at a higher distance difference than regress to Rule I, and (3) a sudden jump if the child remains in Rule II after the jump.

The simulation started with an activation of the distance chunks below the retrieval threshold, resulting in Rule I behavior. To be able to improve to Rule II behavior, the activation of the distance chunk had to increase to reach the threshold. The experimental manipulation of Jansen and Van der Maas is incorporated into the model by assuming an additional saliency related activation component. Hereto, the chunk activation formula presented earlier (Activation = B = $\ln(\sum_{j=1}^{n} t_j^{-d})$) is rewritten to: Activation = B + saliency. We assume the additional saliency activation to be linearly related to the distance difference (saliency = Δdistance/$\underline{c}$, where $\underline{c}$ is a scaling constant).

As transition to Rule II is explained by an increase in activation, regression to Rule I behavior is explained by a decrease in saliency activation. After the item with the most extreme distance difference has been presented, the saliency component of the activation equation decreases. However, during Rule II behavior, the distance feature chunk has been retrieved a number of times. This increases the base activation of the chunk which in some cases is sufficient to remain in Rule II for lower levels of the additional saliency activation.

Given these mechanisms, what causes one child to show behavior according to the Maxwell convention, whereas others behave according to the Delay convention or even show a sudden jump? A possible explanation lies in differences of thoroughness of processing. For example, one child might make a mental "note" after discovering a new method to solve the problem, whereas another child might be less inclined to elaborate the new method. We incorporated this explanation into the model. Elaboration is modeled by extra retrievals of the involved chunks, leading to an increase in the number of references. Differences in elaboration of the new method leads to differences in activation increase.

Simulations have shown that this notion of differences in elaboration was sufficient to simulate all observed transition patterns. The results of these simulations are plotted in Figure 4. Each of the panels of this figure contains the base activation, the additional saliency activation and the combination of these two. If the combination of both activations reaches threshold (the dotted horizontal line at y=0), the model is able to retrieve the distance feature, and performance will improve to Rule II (shown by the dotted line with round markers). Figure 4 shows that the base activation increases for all three panels if the distance feature is retrieved. The size of the increase depends on the amount of elaboration expressed in retrievals. For the left-panel, the Maxwell convention, we assumed that the distance chunk was retrieved only once resulting in one additional reference, the delay convention occurs if the distance chunk is retrieved and elaborated resulting in three additional references, and the sudden jump occurs if retrieving and elaboration lead to five additional references.

Besides explaining the transitions, this mechanism also gives an explanation of the torque

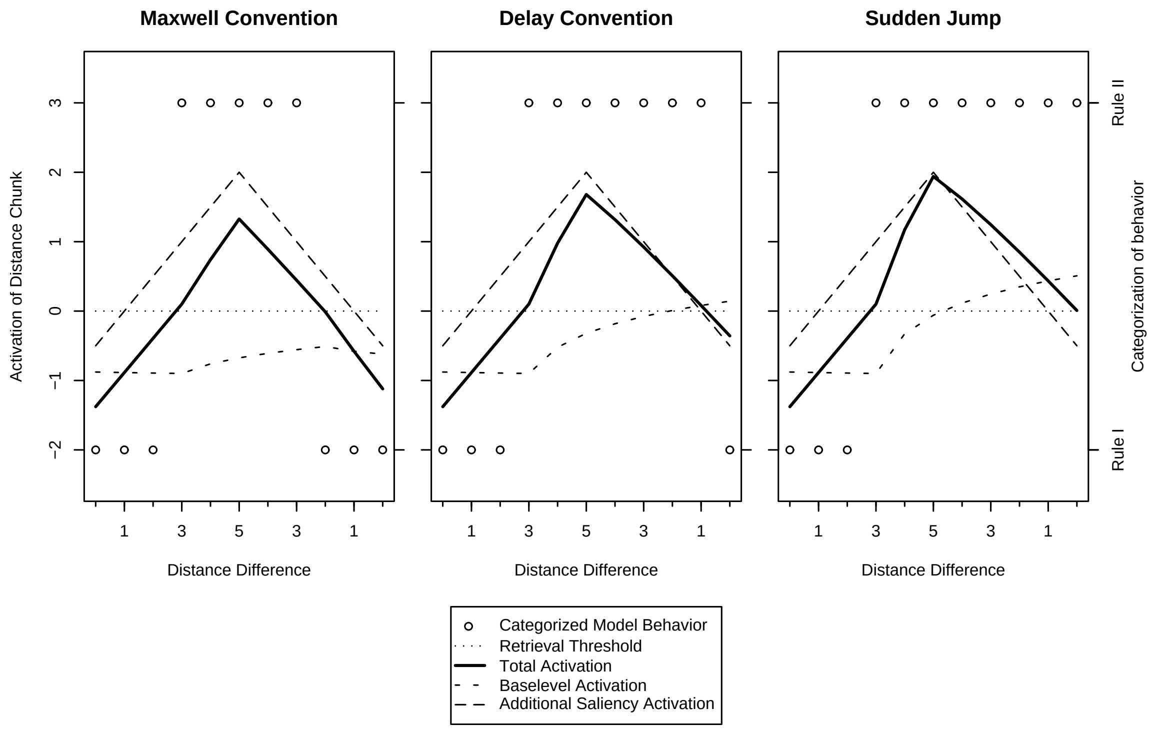

Figure 4. Hysteresis patterns explained by the retrieval of the distance feature.

difference effect. Because torque difference is strongly correlated with distance difference, a large torque difference increases the activation of the chunk that represents the distance dimension. This suggests that the torque difference effect is mainly found during the periods around the transitions. This is in accordance with the results of Jansen and Van der Maas (1997), showing that although a Torque Difference Effect is observable for extreme cases, the homogeneity assumptions still hold.

Discussion of the Model With respect to EC1, the model constructs new production rules. These new production rules are then subjected to expected gain learning. When the expected gain of these new production rules drops below the expected gain of the interpretation proceduralization rules, the model tries to find a better method to solve presented balance scale items. By means of this process, the model is able to learn new production rules, for example for Rule III and Rule IV.

However, because no detailed information is available on the transitions to Rule III and Rule IV, we did not incorporate these transitions in the model. Therefore, this models satisfies EC2 only for the transition from Rule I to Rule II. If more detailed information on the transitions to Rule III and Rule IV comes available, these transitions can be easily incorporated into the model.

The model presented above explains transitions without feedback and also shows the three transition patterns as presented in Figure 2. Therefore, this model reproduces the no-feedback learning phenomenon (EC3). This model also reproduces the torque difference effect. Because this model only exhibits the Torque Difference Effect around transitions and for item types which are better answered in the next, more complex Rule, the model still adheres to the homogeneity of item types. This satisfies EC4.

Summarizing, we showed that a model based on small set assumptions and a principled proceduralization mechanism is able to simulate human development. At the same time, this model is able to simulate and explain individual development patterns. Moreover, the model showed that the torque difference effect (Ferretti & Butterfield, 1986) might be caused by the unstable period around the transition, instead of being a global effect.

Discussion

Four different phenomena can be distinguished in development on the balance scale task. First, behavior of children has been shown to be classifiable into qualitatively distinguishable phases. Second, behavior in these phases can be described by a set of increasingly complex Rules (see Figure 1). Third, recent research has identified particular transition patterns during the transition from Rule I to Rule II. Moreover, no feedback is necessary to induce this transition. Fourth, given a large difference between the torques on left and right side of the balance scale, behavior is likely to be according to a more advanced Rule than if the difference is smaller (the Torque Difference Effect).

Since the first description of the Rules, a number of computational models have been proposed as models of development on the balance scale task. The production rule model of Sage and Langley (1983) reproduces transitions, but the structure of the model led to the construction of production rules that deviated from the empirically identified Rules. Another approach was taken by McClelland and Shultz et al. (McClelland, 1989, 1995; Shultz & Schmidt, 1991; Shultz et al., 1994, 1995). They constructed different types of connectionist models which showed feedback mediated transitions. However, reanalysis of the behavior of these models has shown (Jansen & Van der Maas, 1997) that the behavior of these models is not accurately described by qualitatively different rules. A third approach was taken by Schmidt and Ling (1996). They modeled behavior on the balance scale task with a model that constructs decision trees. With a number of additional assumptions concerning the input representation, their model showed behavior comparable to human data. However, the plausibility of both the internal representation of the Rules (especially the problem-to-answer mappings of Rule IV) and the additional input related assumption is questionable.

Taken together, the two ACT-R models presented here meet the empirical criteria. The first ACT-R model reproduces the correct sequence of Rules. The model advances through the observed Rules by replacing production rules for less succesful Rules with production rules for more advanced Rules. The model makes three key assumptions. The first assumption is that the knowledge used to solve a balance problem varies in costs, representing difficulty. Because of this assumption, the model does not rely on ad hoc distributions of balance scale items. The second assumption of the model is that parameters of the production rules that reflect success and failure are initialized with a number of prior successes and failures. This results in a more stable learning pattern. Without these initial experiences the expected gains of the production rules are heavily influenced by a single successful or unsuccessful application of the production rule. The third assumption of this model is the availability of the production rules. How the rules are constructed is shown for Rule I and Rule II by the second model. The central notion of this model is that a general mechanism guides the discovery of the Rules. This mechanism is associated with the detection of differences. If a balance scale problem is presented, children have to predict to which side the scale tilts. For this, they have to find a difference between the left and right side. Children who use Rule I have discovered the importance of the weight dimension, children using Rule II have also grasped the importance of the distance dimension. This is incorporated into the model by assuming a necessary retrieval of the dimension to be used in the reasoning process. If retrieval is not possible, either because the dimension is not available at all or because its activation is below a retrieval threshold, then the model does not perform at the level for which the retrieval was necessary. If the retrieval succeeds and proves successful by producing correct answers, this newly discovered reasoning process is incorporated by the construction of new production rules which automate the retrieval of the dimension. The new rules prove useful and their costs gradually decrease and so they come to dominate behavior. Because of this, new Rules are learned. The model based on this idea, is shown to learn Rule I and Rule II. Rule III and Rule IV are also learnable by the model, assuming knowledge about the combination of two dimensions is available to the model. However, because the lack of detailed empirical data related to these transitions, these transitions are not incorporated in the second model.

The models reproduce the same phenomena found in cognitive development. First, the models' behavior is classifiable into qualitatively distinguishable phases. Second, the first model shows the correct sequence of Rules. Furthermore, the second model gives a detailed account of the transition from Rule I to Rule II. Third, the transition model shows development without being dependent on feedback. We assume that because of visual saliency, the distance dimension is more likely to be noticed given larger differences between the left and right side of the balance. This is reflected in the model by the activation of the memory chunk that represents the distance dimension. Given a larger difference between the distances on the left and right side of the balance, the activation of the distance dimension chunk might increase to just above the retrieval threshold. If this occurs, the behavior of the model shifts from Rule I to Rule II. This mechanism gives an explanation for the development from Rule I to Rule II without feedback. This idea also accounts for the transition patterns observed by Jansen and Van der Maas (in press). If an increase in difference is able to push the activation over the threshold, a decrease in difference might cause a drop below the threshold. Combined with an increase in activation caused by the retrievals, the identified transition patterns emerge. Fourth, because larger distance differences are related to larger torque differences, this model shows the Torque Difference Effect during the transition periods. If the distance differences increase, the chances of applying Rule II increase. The same effect as implied by the Torque Difference Effect. However, the model predicts this effect to occur during transitions only. When behavior is stable, no chunks need to be retrieved, so no influence of torque differences is observed in the model. During stable behavior, the model will show homogeneous answers for problems of the same item type. This is in line with the conclusion that the Torque Difference Effect is solely observable for large toque differences (Jansen & Van der Maas, 1997).

Summarizing, we have shown how a model that is capable to show the standard phenomena of qualitative Rules and observed order of Rules. To model these phenomena, the model is constructed on a basis of clearly stated assumptions. With respect to the learning without feedback and the transition patterns, we have presented a model that is able to improve from Rule I to Rule II without feedback. As far as we know, this is the first model able to model development in the balance scale domain without being dependent on feedback. Furthermore, this model is able to show the transition patterns related to this transition. This same model is also able to explain the Torque Difference Effect. However, instead of presenting the Torque Difference Effect as a global effect, this model suggests that the Torque Difference Effect is related to the less stable periods around a transitions.

References

Anderson, J. R. (1987). Methodologies for studying human knowledge. Behavioral and Brain Sciences, 10(3), 467–505.

Anderson, J. R. (1990). The adaptive character of thought. Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Anderson, J. R., & Lebiere, C. (1998). The atomic components of thought. Mahwah, NJ: Lawrence Erlbaum Associates, Inc.

Boden, M. A. (1979). Piaget. Glasgow, UK: William Collins Sons & Co. Ltd.

Ferretti, R. P., & Butterfield, E. C. (1986). Are children's rule-assessment classifications invariant across instances of problem types? Child Development, 57(6), 1419–1428.

Ferretti, R. P., & Butterfield, E. C. (1992). Intelligence-related differences in the learning, maintenance, and transfer of problem-solving strategies. Intelligence, 16(2), 207–223.

Ferretti, R. P., Butterfield, E. C., Cah, A., & Kerkman, D. (1985). The classification of children's knowledge: Development on the balance-scale and inclined-plane tasks. Journal of Experimental Child Psychology, 39(1), 131–160.

Jansen, B. R. J., & Van der Maas, H. L. J. (1997). Statistical test of the rule assessment methodology by latent class analysis. Developmental Review, 17, 321–357.

Jansen, B. R. J., & Van der Maas, H. L. J. (in press). Evidence for the phase transition from Rule I to Rule II on the balance scale task. Developmental Review.

Klahr, D., & Siegler, R. (1978). The representation of children's knowledge. In H. Reese & L. Lipsitt (Eds.), Advances in child development and behavior (pp. 61–116). New York, NY: Academic Press.

Langley, P. (1987). A general theory of discrimination learning. In D. Klahr, P. Langley, & R. Neches (Eds.), Production system models of learning and development. Cambridge, MA: MIT Press.

Lebiere, C., & Anderson, J. R. (1998). Cognitive arithmetic. In J. R. Anderson & C. Lebiere (Eds.), The atomic components of thought. (pp. 297–342). Mahwah, NJ: Lawrence Erlbaum Associates, Inc.

Lovett, M. C., & Anderson, J. R. (1996). History of success and current context in problem solving. Cognitive Psychology, 31(2), 168–217.

McClelland, J. L. (1989). Parallel distributed processing: Implications for cognition and development. In M. G. M. Morris (Ed.), Parallel distributed processing, implications for psychology and neurobiology (pp. 8–45). Oxford: Clarendon.

McClelland, J. L. (1995). A connectionist perspective on knowledge and development. In T. J. Simon & G. S. Halford (Eds.), Developing cognitive competence: New approaches to process modeling (pp. 157–204). Lawrence Erlbaum Associates.

Newell, A. (1990). Unified theories of cognition. Cambridge, MA: Harvard University Press.

Normandeau, S., Larivee, S., Roulin, J. L., & Longeot, F. (1989). The balance-scale dilemma: Either the subject or the experimenter muddles through. Journal of Genetic Psychology, 150(3), 237–250.

Raijmakers, M. E. J., Van Koten, S., & Molenaar, P. C. M. (1996). On the validity of simulating stagewise development by means of PDP networks: Application of catastrophe analysis and an experimental test of rule-like network performance. Cognitive Science, 20(1), 101–136.

Reese, H. W. (1989). Rules and rule-governance: Cognitive and behavioristic views. In

S. C. Hayes (Ed.), Rule governed behavior: Cognition, contingencies, and instructional control. (pp. 3–84). New York: Plenum.

Sage, S., & Langley, P. (1983). Modeling cognitive development on the balance scale task. In A. Bundy (Ed.), Proceedings of the eight international joint conference on artificial intelligence (Vol. 1, pp. 94–96). Karlsruhe, Germany.

Schmidt, W. C., & Ling, C. X. (1996). A decision-tree model of balance scale development. Machine Learning, 24, 203–230.

Shultz, T. R., Mareschal, D., & Schmidt, W. C. (1994). Modeling cognitive development on balance scale phenomena. Machine Learning, 16, 57–86.

Shultz, T. R., & Schmidt, W. C. (1991). A cascade-correlation model of balance scale phenomena. In Proceedings of the thirteenth annual conference of the cognitive science society (pp. 635–640). Hillsdale, NJ: Erlbaum.

Shultz, T. R., Schmidt, W. C., Buckingham, D., & Mareschal, D. (1995). Modeling cognitive development with a generative connectionist algorithm. In T. J. Simon & G. S. Halford (Eds.), Developing cognitive competence: New approaches to process modeling (pp. 205–261). Lawrence Erlbaum Associates.

Siegler, R. S. (1976). Three aspects of cognitive development. Cognitive Psychology, 8, 481–520.

Siegler, R. S. (1981). Developmental sequences within and between concepts. Monographs of the Society for Research in Child Development, 46(2), 1–74. (With commentary and reply from Sidney Strauss and Iris Levin)

Siegler, R. S. (1989). Mechanisms of cognitive change. Annual Review of Psychology, 40, 353–379.

Siegler, R. S., & Chen, Z. (1998). Developmental differences in rule learning: A microgenetic analysis. Cogntive Psychology, 36(3), 273–310.

Simon, T. J., & Halford, G. S. (Eds.). (1995). Developing cognitive competence: New approaches to process modeling. Lawrence Erlbaum Associates.

Taatgen, N. (2000). Proceduralization in ACT-R (Tech. Rep.). Groningen University, The Netherlands. (Retrieved June 13th, 2001, from World Wide Web: `http://tcw2.ppsw.rug.nl/~niels/proceduralization/proceduralization.doc`.)

Taatgen, N. A., & Anderson, J. R. (submitted). Why do children learn to say broke? Cognition.

Thom, R. (1975). Structural stability and morphogenesis. New York: Benjamin Addison Wesley.