In: Proceedings of the European Conference on Cognitive Science, 1997.

MODELLING TRANSITIONS IN CHILDREN'S DEVELOPMENT BY STARTING WITH ADULTS

Gary Jones & Frank E. Ritter
Psychology Department
University of Nottingham
Nottingham NG7 2RD
England
E-mail: gaj@psychology.nottingham.ac.uk
ritter@psychology.nottingham.ac.uk

17th March, 1997

ABSTRACT

We present a blocks task in which cognitive change is apparent, and propose that computational modelling can help in examining how cognitive change occurs. We initially examine the performance of adults on the task, and use the results of this experiment to build both an adult model and a simulation of the task. Comparisons of the model with the adult data reveal further completion work regarding the models qualitative behaviour. However, when finished, the model will be able to predict how behaviour changes when domain knowledge is removed, and will provide the basis for models of children's behaviour on the task.

Keywords: Cognitive modelling, child development, cognitive change, interaction, simulation.

Acknowledgements
The Office of Naval Research provided initial support for this work under contract N00014-96-M-0012, as did the Society for Artificial Intelligence and the Simulation of Behaviour.

INTRODUCTION

The need to look at transitions in cognitive development has been known for quite some time (e.g. Simon, 1962). However, most literature in children's development has tended to concentrate on describing children's behaviour at each performance level with little regard to how progression from one level to another occurs (e.g. Piaget & Inhelder, 1969). Siegler and Shipley (1995) see this as being problematic because their studies suggest that it is misleading to view children at one age as thinking in one specific way. They also believe that if children's development is thought of in terms of these one-to-one correspondences, then this presents large gaps for theories of transition to fill. They therefore put forward the need to examine transitions in conjunction with performance.

The solving of physical problem solving puzzles is a good area in which to examine both children's performance and the transitions therein. This is because a detailed analysis of the task behaviour is possible via videotape. Many strategies will therefore be readily visible, reducing the need for the experimenter to infer what mental strategies are being used. For this reason we use a physical problem solving puzzle, the "Tower of Nottingham", to study transitions in children's performance and the factors influencing them.

THE TASK

The Tower of Nottingham is a puzzle which consists of 21 blocks, with a goal to produce a pyramid structure (see Figure 1). There are six layers to the pyramid, the lower five consisting of four blocks each, with a single block as the top layer. The blocks in the lower five layers all share the same characteristics (as shown in Figure 2), differing only by size. Two of the blocks have half-pegs, with one block having a hole and the other a peg, such that placing the peg in the hole brings the half-pegs together to form a peg. Similarly, the other two blocks have half-holes, and placing the peg of one block into the hole of the other produces a hole. A square shaped layer is then produced by inserting the newly created peg into the newly created hole. The block features also permit the construction of a layer in a variety of other ways.

Two further features exist: each block has a quarter circle indent on top and a quarter circle depression underneath. When a layer is created, the quarter circles form circles in the centre such that layers can be stacked on top of each other by placing the circular depression of the upper layer onto the circular indentation of the lower layer.

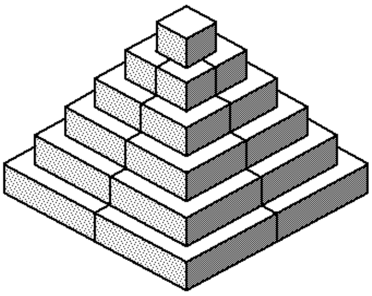

Figure 1: The final assembly of the 21 blocks that comprise the Tower of Nottingham.

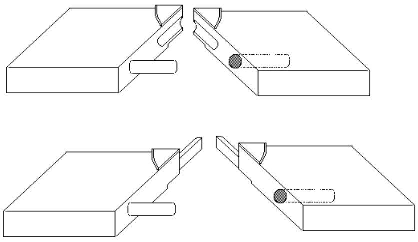

Figure 2: The four blocks that make up each of the lower five layers in the Tower of Nottingham.
(Figures courtesy of Peter Cheng & Heather Wood).

In studies involving this task (e.g. Wood & Middleton, 1975) children show a progression in performance with age, such as a reduction in errors and time taken, and an increase in the correct operations accomplished. Older children are also more efficient at the task, such that they no longer use some of the inefficient strategies that younger children use (e.g. no longer sampling with replacement).

All of the studies performed so far have involved instruction, and the performance progression is also seen with regard to how much reliance the children placed on instructional aids. For example, the tutor had to intervene more often with younger children (Wood, Bruner & Ross, 1976) and younger children are less successful in utilising nine pictures showing the stages of completion of the puzzle (Murphy & Wood, 1981).

EXAMINING TRANSITIONS

Since there are different performance levels in this task, transitions must occur in order to transcend these performance levels. Whilst it is relatively simple to characterise behaviour at each performance level, defining how, why and what transitions take place is problematic. Siegler (1995) puts forward the microgenetic approach as a method by which transitions may be studied. In this approach, behaviour is observed in detail and as much as possible during the period of transition. However, we see two problems with this approach. First, the exact period when transition takes place is likely to be different across individuals. Second, the data by itself is not sufficient to prove that a defined transition mechanism can cause the necessary changes in behaviour.

Computational modelling is one method that can supplement the microgenetic approach and help explore theories of transition. Models enable learning and knowledge to be independently and directly manipulated. This means that models can test what initial knowledge is required to produce the behaviour seen at each performance level, and predict how transitions and different levels of behaviour may occur. Such predictions can be tested against the thorough analysis of change that the microgenetic approach provides.

The use of computational modelling requires defining the behaviour that occurs at each performance level, since the model requires the knowledge and procedures that children may be using. To the extent that the behaviour cannot be defined, the model can make predictions as to what the missing elements could be. Therefore modelling task behaviour can provide a method for examining to what extent changes in task performance can be attributed to differences in knowledge and to what extent changes in task performance can be attributed to developmental processes.

The modelling approach has only been used to a limited extent in developmental psychology. McClelland and Jenkins (1991) produced a model of the balance-beam task which compared favourably to the performance levels of children found by Siegler (1981). In their model, the transition mechanism is essentially driven by experience with the task, and not development per se. Jones and VanLehn (1991) detail a model which accounts for the sum-to-min strategy transition whereby transition occurs by seeking to make strategies more efficient (though this ignores other task strategies that take place, such as retrieval). Although these two models are relatively simple, they show that computational modelling can indeed provide accounts of the possible mechanisms that occur in cognitive change.

There would appear to be two approaches with regard to creating a model in our task domain. One method is to model a lower performance level and see if that model can then progress to the higher performance levels that we see on the task. The other method is to begin at the highest performance level (that of adults), and then see if reduced versions of this model show behaviour that looks like lower performance levels. In both cases, the model should make clear the predictions it makes, both in terms of hypotheses as to what the missing elements are, and in terms of task predictions that have not been examined yet (this is a test for any computational model).

The problem with beginning at the lower performance levels is that young children on the Tower of Nottingham often generate complex behaviour due to lack of knowledge. It can therefore be quite difficult to ascertain what strategies and what initial knowledge they may have. On the other hand, adults on the task can be considered as being at the highest performance level, and are also able to give verbal protocols. This will help to provide a clearer picture on what strategies and knowledge they have when starting the task. We have therefore examined adult behaviour in the Tower of Nottingham, in order to provide a baseline for the level of performance that children will eventually attain.

ADULT BEHAVIOUR ON THE TASK

Ten adults attempted to build the Tower of Nottingham whilst giving verbal protocols. In stage one, half of these were shown a picture of the completed tower prior to beginning the experiment (the “goal” condition), with the other half being told to “build something special” (the “non-goal” condition). Once the tower had been completed, it was dismantled by the experimenter (out of view of the subject) and subjects were asked to re-build it (stage two). All subjects managed to complete the tower in both stages.

As can be seen in Table 1, subjects in the goal condition were faster and produced less errors than their non-goal counterparts in stage one of the experiment. This advantage almost disappears in stage two. All subjects completed the task quickest when doing it for the second time, and for all subjects there was a trend towards taking less time per layer the further into the task they were, in excess of the reduced search time required.

	STAGE ONE		**STAGE TWO**	
	Goal	**Non-Goal**	**Goal**	**Non-Goal**
Time taken	149	319	101	99
Errors selecting blocks	7	13	7	3
Errors fitting blocks	57	89	32	34

Table 1: Timings (in seconds) and errors for experiment stages one and two.

These results provide evidence that some form of learning is occurring throughout the task for all subjects. What we can ascertain from video and protocol analysis is that adults start with the basic knowledge that is required to complete the task, such as pegs can go into holes, half-pegs can fit to other half-pegs to make pegs, pyramids are made from items of different sizes, four quarter circles make a circle, and so on. What subjects must be learning, therefore, is how to apply this knowledge to the task at hand. In stage one, the goal subjects must first learn how to build layers, and the non-goal subjects must learn both this and that the blocks make a pyramid. This can be seen as the main reason for subjects being faster in stage two. The within stage learning that is evident occurs after this new knowledge is in place, and therefore must come from experience with the task.

An examination of the strategies used reveals a variety of strategy use in both the selection of blocks and in the application of those blocks to produce layers. Block selection is by size 50% of the time, and features 30%. Subjects produce layers in a variety of ways, the most common methods being via a peg pair and hole pair (40%), and via a pair-threeblock-layer configuration (38%). The remaining strategy is to create a layer using a pair having two pegs and a pair having two holes.

THE MODEL

A preliminary model of the task has been created in ACT-R (Anderson, 1993), which interacts with a graphical simulation of the tower in Garnet (Myers et al, 1990). We will explain the simulation first as it provides input to the model.

The simulation environment includes all of the blocks in a graphical display window provided by Garnet. Presently this is only two-dimensional, so the simulation can only rotate blocks in the x-y plane, though blocks can be turned upside down. This is not a great limitation since we are able to capture most of the task behaviour without simulating three dimensions. All block features are represented, as well as the block and feature dimensions, the orientation of the block, and any other blocks it is connected to.

The simulation includes an eye, whereby information in the fovea, parafovea and periphery is represented. The eye is able to attend to any block and returns all visible features of blocks that are in the fovea and parafovea (the latter is subject to some noise). Neither features nor size is given for peripheral blocks. The time taken to complete one eye movement together with one fixation is 250 ms. This is based on timing estimates detailed by Baxter and Ritter (1996) for saccades and fixations on non-complex stimuli, but does not take into account the distance that the eye is required to move.

We also include a simple hand simulation, which is currently able to pick up, drop, rotate and fit together blocks. Movements of the hand take approximately 550 ms based on estimates from the video data.

The model is based on the ACT-R theory of activation. All declarative memory structures have an associated activation level. Each declarative memory structure that shares similar characteristics to the current goal has its activation altered based on a numeric weighting which indicates whether the structure influences the goal in a positive or negative way. The activation of memory structures is important because when several productions can fire, the one with the highest activation is selected. The total activation of a production is the sum of the activations for all structures in its condition (which is the basis for calculating production latency).

Our model holds all productions as being equal such that selection between productions is based on activation. By manipulating the numeric weighting between the goal and declarative memory structures, we can in effect alter the activation of declarative memory structures and thus can influence which production will fire.

The method by which we manipulate the influence that declarative memory structures have upon the goal can be viewed as being similar to the contingent method of tutoring (e.g. Wood & Middleton, 1975), which appears to be a rational explanation of the data. When there is a block in each hand, there are many ways to fit them together. Since activations and numeric weights are all equal when beginning the task (actually, some features could be more salient and so should possibly have a higher activation, but this is not currently implemented), the production to fire is selected at random. If this results in a construction which we think may be correct, we alter nothing. However, if we think the construction is incorrect, we decrease the likelihood of blocks being put together using that method again by decreasing the numeric weighting between the goal and memory structures. Note that activation levels are subjected to a small amount of noise.

Each production that fires takes 50 ms (the default production latency in the ACT-R architecture). Added to this is the time taken for each of the declarative memory structures in the condition of the production to be matched. This is manipulated via the numeric weighting, but we begin with all structures having an activation of 1.0.

The model begins with the initial task knowledge that subjects have, based upon the video analyses (and detailed earlier). They know that pegs can go into holes, half-pegs can go into half-holes, and that quarter circles can make a circle. In addition, they know that two half-pegs can make a peg, and two half-holes can make a hole. Assessment of a “good” construction is mainly based on whether it has straight outer edges (this criteria was used by most subjects). Other criteria are quarter circle alignment, whether the features fit well, and whether the construction lies flat on the table. These are all consistent with the behaviour of adults on the task.

The model proceeds by first asking the eye to look at the blocks on the table. If the model has the knowledge that blocks of the same size go together, then it attempts to retrieve blocks of the same size. If the model has no knowledge about blocks of the same size, it either seeks any two blocks, or analyses what is on the table to try and hypothesise what the blocks make. Once two blocks are selected, the model must decide how to fit them together. This decision is made based upon what features each block has, and what knowledge the model has of those features. Once fitted, the construction must be assessed and a decision to either continue or disassemble is made. The block selection process then continues. This is all accomplished via a high degree of interaction between the model and the simulation.

COMPARISON OF THE MODEL WITH THE DATA

The model was run five times, and included the fact that blocks of the same size go together as part of its initial knowledge (i.e. the goal condition). As we can see from Table 2, the timing data show that the model is performing in close accordance with the goal subjects in stage one of the experiment, with a correlation of 0.87 across levels.

	Size 6	**Size 5**	**Size 4**	**Size 3**	**Size 2**	**Total**
MODEL	31.3	25.3	23.7	23.8	19.4	123.5
SUBJECTS	34.0	28.4	21.4	15.0	16.0	114.8

Table 2: Timings (in seconds) for time spent on each layer (subject times are from goal condition, stage one).

A qualitative analysis of the model’s behaviour is less conclusive. Selection of blocks is always by size only, and the strategy used is always pair-threeblock-layer (this is a current limitation of the simulation). Although blocks in the parafovea are subjected to noise, there were no errors in selection by size in all 5 runs of the model. Errors for fitting blocks were 46 as compared to the 57 that the goal subjects made in stage one.

SUMMARY

The comparison of the subject data with the model reveals that the simulation and model require some updating in order to show behaviour which qualitatively fits the subject data. Once complete, the model will enable distinct predictions to be made regarding performance when task characteristics are altered, such as how well perception works (which is an area we expect true development to occur). It should also make predictions regarding task difficulty and time to complete the task when subjects are given different amounts of task knowledge. Both of these may well influence cognitive development on this task.

The first area that we wish to investigate is how behaviour changes when domain knowledge is removed (either declarative knowledge, procedural knowledge, or both). This should help determine to what extent children’s behaviour may be due to lack of knowledge and to what extent it may be due to developmental influences. This will involve changes to the model’s architecture and the eye simulation. Further work will then involve producing models which fit the behaviour of 7, 5 and 3 year olds on the tower, and then analysing these to see how transitions between them may give rise to the advanced behaviours observed.

REFERENCES

Anderson, J.R. (1993). Rules of the mind. Hillsdale, NJ: Lawrence Erlbaum Associates.

Baxter, G.D. & Ritter, F.E. (1996). A simple set of generic perceptual and motor capabilities for use by cognitive models in SLGMS. *CREDIT Technical Report No.36, Psychology Department, University of Nottingham* (available from Frank E. Ritter, on request).

Jones, R.M. & VanLehn, K. (1991). Strategy shifts without impasses: A computational model of the sum-to-min transition. In *Proceedings of the Thirteenth Annual Conference of the Cognitive Science Society, 358-363.*

McClelland, J.L. & Jenkins, E. (1991). Nature, nurture and connections: Implications of connectionist models for cognitive development. In VanLehn, K. (Ed.) *Architectures for intelligence, Chapter 3.*

Murphy, C.M. & Wood, D.J. (1981). Learning from pictures: The use of pictorial information by young children. *Journal of Experimental Child Psychology, 32, 279-297.*

Myers, B.A., Guise, D.A., Dannenberg, R.B., Vander Zanden, B., Kosbie, D.S., Pervin, E., Mickish, A. & Marchal, P. (1990). Comprehensive Support for graphical, highly-interactive user interfaces. *IEEE Computer, 23 (11), 71-85.*

Piaget, J. & Inhelder, B. (1969). The psychology of the child. New York: Basic Books.

Siegler, R.S. (1981). Developmental sequences within and between concepts. *Monographs of the Society for Research in Child Development, 46 (No. 189, 1-74).*

Siegler, R.S. (1995). Children's thinking: How does change occur? In Schneider, W. & Weinert, F. (Eds.), *Memory performance and competencies: Issues in growth and development.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Siegler, R.S. & Shipley, C. (1995). Variation, selection and cognitive change. In Simon, T. & Halford, G. (Eds.), *Developing cognitive competence: New approaches to process modeling.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Simon, H.A. (1962). An information processing theory of intellectual development. *Monographs of the Society for Research in Child Development, 27 (2, Serial No. 82).*

Wood, D.J., Bruner, J.S. & Ross, G. (1976). The role of tutoring in problem solving. *Journal of Child Psychology and Psychiatry, 17, 89-100.*

Wood, D.J. & Middleton, D. (1975). A study of assisted problem solving. *British Journal of Psychology, 66(2), 181-191.*