

Jones, G., Ritter, F. E., & Wood, D. J. (2000). Using a cognitive architecture to examine what develops. *Psychological Science*, 11, 93-100.

RUNNING HEAD: EXAMINING WHAT DEVELOPS

Word Count: 4991

Using a Cognitive Architecture to Examine What Develops

Gary Jones, Frank E. Ritter, David J. Wood

ESRC Centre for Research in Development, Instruction and Training

School of Psychology, University of Nottingham, Nottingham NG7 2RD (UK)

Send correspondence about this article to:

Gary Jones
School of Psychology
U. of Nottingham
Nottingham NG7 2RD
UK
Email: gaj@psychology.nottingham.ac.uk
Phone: +44 (115) 951-5151 x18348
Fax: +44 (115) 951-5324

Abstract

Different theories of development propose alternative mechanisms by which development occurs. Cognitive architectures can be used to examine the influence of each proposed mechanism of development whilst keeping all other mechanisms constant. An ACT-R computational model was created which matched adult behavior in solving a 21 block pyramid puzzle. The model was modified in three ways, each of which corresponded to mechanisms of development proposed by developmental theories. The results showed that all the modifications (two of capacity and one of strategy choice) could approximate the behavior of seven year old children on the task. The strategy choice modification provided the closest match on the two central measures of task behavior (time taken per layer, $r=0.99$, and construction attempts per layer, $r=0.73$). Modifying cognitive architectures is a fruitful way to compare and test potential developmental mechanisms, and can therefore help in specifying "What develops?".

Introduction

Examining how children's performance improves, and detailing what changes occur during development, is difficult. This is mainly because development occurs over a number of years. Observations spanning several years are necessarily sparse which means that the mechanisms giving rise to the changes in performance have to be inferred (Siegler, 1995). Many theories of development are based on these sparse observations, so the developmental mechanisms they propose suffer problems of being ill-defined and hence difficult to test.

These problems can be addressed at a micro-level by using the microgenetic method. This generally takes the form of intense study of task behavior over a period of weeks when cognitive change occurs within the task (Siegler & Jenkins, 1989). The microgenetic method offers valuable insights into cognitive change because of the extensive amount of data collected. The limitation of the method is that it covers only a small period of time – development occurs over a number of years.

Computational modeling can help to both force precision and enable the testing of proposed theoretical mechanisms. Modeling enables the manipulation of single variables whilst reliably controlling for all others. A model can therefore be used to test the influence of one theoretical mechanism of development whilst keeping all other mechanisms constant. The program-like nature of computational models also forces the complete definition of any ill-defined concepts that are part of a theoretical mechanism of development.

Computational models have a long but not extensive history within development. Young (1976) and Klahr and Siegler (1978) have modeled improvement in task behavior by increasing the knowledge base of their models. McClelland and Jenkins (1991) modeled development by repeatedly training their model on task problems, corresponding to more task experience. Shultz, Schmidt, Buckingham and Mareschal (1995) modeled development by giving the model more experience on the task, whilst allowing the model to alter its architecture.

Each of these models matched the data from children in their respective tasks. However, problems remain. First, development is believed to be more than simple increases in the knowledge base (Gentner, Rattermann, Markman & Kotovsky, 1995). Second, development is not simply task

experience (Elman, Bates, Johnson, Karmiloff-Smith, Parisi & Plunkett, 1996). Third, the models only match subject data on a small number of measures, with every model omitting timing data. Timing data includes both within-task times (the time taken to perform different components of a task or problem) and solution times (the time taken to perform the whole task or problem). No models have examined within-task times, although relative solution times have been examined by Siegler and Shipley (1995), and Siegler and Shrager (1984). Timing data are important because they can help highlight where learning takes place, and provide objective and precise comparisons between model and subject data.

The central mechanisms of development that have been proposed are knowledge changes¹ (e.g., Klahr & Wallace, 1976; Piaget, 1950), increased capacity (Case, 1985; Halford, 1993; Pascual-Leone, 1969), strategy choice (Siegler & Shipley, 1995), and processing speed (Kail, 1991). Proponents of each mechanism generally admit that there are several mechanisms at work; the mechanisms interact with each other to mediate development. For example, knowledge changes may be influenced by capacity (Case, 1985).

This paper presents the first steps towards testing developmental mechanisms within a computational framework. Examining the influence of each mechanism of development within a single task will help clarify the roles each play in cognitive development. The influence of each mechanism of development within a single task can be tested independently within a computational model of the task. A task to study development is outlined. A description of a computational model of the task, together with its fit to the behavior of adult subjects on the task, is detailed. Three example modifications to the model are presented, which independently manipulate two mechanisms of development: capacity and strategy choice (Jones, 1998, examines all four of the mechanisms listed above). The discussion examines the results of the modifications and presents a program for future research.

A Task to Study Development

The Tower task (Wood & Middleton, 1975) is to build a pyramid from 21 wooden blocks. Figure 1 shows the six layers to the pyramid; the lower five consist of four blocks, with a single

block as the top layer. The blocks in the lower layers share the same characteristics, differing only in size. Each layer is normally formed via two sets of paired blocks. For example, placing the peg of block A into the hole of block B brings the two half-holes together to form a pair having a hole (a hole-pair). Similarly, placing block C and block D together forms a pair with a peg (a peg-pair). A layer is then formed by fitting the peg of the peg-pair into the hole of the hole-pair. The layers have interlocking circular indents above and circular depressions beneath. Layers can thus be stacked to form the Tower. The optimal number of constructions required to build the Tower is twenty, three for each of the lower five layers, and five to stack layers.

Typical errors include fitting half-pegs into holes, placing blocks flush to each other without having a peg/hole connection between them, and putting blocks of different sizes together. Incorrect constructions are rarely left uncorrected, because mismatches are usually readily apparent. It is natural to complete the Tower layer by layer (all subjects in our experiments do this). Within-task learning can occur because the lower five layers all share the same block characteristics.

Insert Figure 1 About Here

Why use this task?

The Tower task is ideal for studying development because it illustrates how children's problem solving behavior changes across ages. Older children accomplish more correct constructions, produce less incorrect constructions, and take less time than their younger counterparts (Murphy & Wood, 1981, 1982; Wood, Bruner & Ross, 1976; Wood & Middleton, 1975). The physical nature of the task enables detailed timing data to be recorded, and allows many strategies to be readily visible, reducing the need for the experimenter to infer mental structures and strategies.

Use of cognitive architectures

Cognitive architectures support the development of computational models within a common theoretical framework. The architectures usually incorporate some general psychological

mechanisms, including those proposed by developmental researchers (e.g., the representation and characteristics of working memory) that constrain the way behavior can be modeled (e.g., Newell, 1990, Chapter 1).

Whilst most current models of development have been successful in being able to produce behavior at different levels of performance by only adjusting knowledge, this does not reveal the extent to which other mechanisms may influence development. Cognitive architectures provide a principled way to examine the extent to which mechanisms other than knowledge may influence development.

In addition, most developmental models have not been created within a cognitive architecture, so they could be manipulated in an unconstrained way in order to fit the subject data. Creating models within cognitive architectures forces constraints on any model developed within the architecture. In particular, the timing predictions that architectures provide enable a direct matching of the temporal behavior of the model with that of subjects.

Simulating Development

There are two natural ways to create a series of developmental models. One way is to model a lower performance level (that of children) and enhance the model to fit higher performance levels. The other way is to begin at the highest performance level (that of adults), and then degrade the model to fit lower performance levels. Adult performance on the task is simpler to model because adult behavior is more regular, and adults can readily provide verbal protocols which help in the creation of a model. A model of adult behavior on the task is therefore created first, and modifications to this model will be compared with the behavior of children on the task.

The task simulation and the model of adult behavior

Both a task simulation and a model are necessary to study development in this task. The task simulation and model are detailed enough to allow modifications to be made which represent the mechanisms proposed by developmental theories. A task simulation is required so that the task is appropriately represented and the time spent interacting with the blocks (i.e., looking at and

manipulating the blocks) is properly estimated. Figure 2 shows the task simulation, which contains a full graphical representation of the task including all blocks and features.

Insert Figure 2 About Here

The simulation includes an eye and two hands. The eye and hands are designed to meet a set of requirements identified for creating a psychologically plausible architecture for interacting with multiple external tasks (Baxter & Ritter, 1996). The architecture has been implemented several times in a variety of user interface management systems and graphics programming languages (Ritter, Baxter, Jones & Young, 1999).

The eye is able to saccade and fixate, and passes to the model what blocks and constructions it sees (e.g., a peg-pair will be represented as a construction of two blocks flush on their outer edges with their quarter circles and half-pegs aligned). The quality of the visual information passed to the model depends upon where blocks are positioned in relation to the simulated eye. Three areas of decreasing visual quality are defined: fovea, parafovea, and periphery. The simulated hands are able to pick up, drop, rotate, turn over, fit, and disassemble blocks.

The cognitive aspects of the model are based on the ACT-R cognitive architecture (Anderson, 1993). ACT-R is a production system based architecture in which knowledge is represented as facts and rules. A rule has two sides, a condition and an action. When the conditions of a rule are met, its action can be taken (i.e., the rule can fire). Each condition normally corresponds to a fact that must be present within the system. Facts in ACT-R are called memory elements, with each memory element consisting of a number of attribute-value pairs.

An example rule from the ACT-R model is shown in Table 1. The rule moves the left hand to a construction as a pre-cursor to grabbing the construction. In order for the rule to fire, the system's current goal must be to grab a construction using the left hand. The left hand must not be near the construction, and the model's eye must be looking at the construction (so that the model knows where to move the left hand). If all these conditions are met, the hand is moved to the construction.

Insert Table 1 About Here

The adult model contains 317 ACT-R rules including task knowledge and knowledge about how to direct the eye and the hands. Within ACT-R, all rules have an associated strength. When several rules apply to the same task situation (i.e., there are several rules whose conditions are satisfied), the rule with the highest strength is selected. The strength of rules is altered based on their success and failure. If a rule helped generate a construction which the model believes to be correct, then the rule's strength is increased. If a rule helped generate a construction which the model believes to be incorrect, then the rule's strength is decreased. The manipulation of the strength of rules means that as the task progresses, the strategies (i.e., rules) which result in success are more likely to be used.

The model begins with general rules for fitting blocks together. When the model has produced a construction that it believes to be a success, a new rule is learned that describes the specific block features that were fit together. The rule is therefore more specific than the previous, general, rule. As the task progresses, the number of rules that can be applied to fit blocks increases, and therefore more emphasis is placed on the mechanism for learning success/failure.

The representations of the blocks and block features in the model all have an associated activation. When the memory elements for the blocks and block features are created, they are set to a base level activation. Activation is then subject to decay on each cycle (a cycle involves the selection and firing of a single rule), but is also raised based on the current goals of the model. When the activation of blocks and block features fall below a specified level (the retrieval threshold) they can no longer be matched in rule conditions. This decay represents a capacity limitation within the model.

Comparison of the model to adult data

Adult subjects (N=5) who had never encountered the task before were shown a picture of the completed Tower. The subjects were then presented with the 21 blocks and asked to construct the

Tower, once only and unaided, whilst giving verbal protocols. Construction behavior was video recorded, together with onset timings to the nearest second.

For the model's timing data, the ACT-R default timing of 50 ms per rule firing is used. For rules involving eye movements (saccades) this timing is consistent, but is increased to 200 ms for rules involving fixations, based on a summary of the vision literature provided by Baxter and Ritter (1996). For rules involving motor actions (fitting, rotating, moving, and disassembling blocks) the timing is increased to 550 ms (Jones & Ritter, 1997). Thus, the time predictions made by the model are absolute predictions; they are not scaled or otherwise adjusted in the analyses and graphs reported below.

The adult model begins with the initial knowledge of the task that subjects had, such as blocks of the same size go together, pegs go in holes, intended constructions are flush on their outer edges, and so on. The initial knowledge is apparent in the construction behavior of adults (e.g., they produced a total of 33 peg-pair or hole-pair constructions, 13 two-peg or two-hole constructions, and never considered producing constructions that were not flush on their outer edges).

The ACT-R expected gain noise parameter adds a random amount of noise to the strength of rules. It was set to 0.04 to randomize the selection of rules having the same strength. The usual settings range from 0.0 (the default) to 1.0 (J. R. Anderson, personal communication, February 4, 1999). The initial activation level given to memory elements also differs from the default, because of the decay mechanism (see Jones & Ritter, 1998, for an explanation). All other ACT-R parameters were either unused or set to the suggested default setting.

In addition to expected gain noise, the model contains other random elements (e.g., when searching for blocks, an unseen block in the periphery is randomly selected). To provide a better approximation to the expected model value, the aggregate data from ten runs of the model are compared to the five adult subjects.

There are at least nine measures of behavior that can be taken for the Tower task, and each can be taken as an overall measure and on a layer-by-layer basis. Overall measures provide a guide to performance on the task (e.g., the time taken to complete the Tower), and layer-by-layer measures show any learning that occurs (e.g., the time to produce each layer). The main measures reported here

are the time taken to complete the Tower, and the number of constructions made in completing the Tower (the number of constructions made refers to the raw number of correct and incorrect constructions produced whilst completing the Tower).

For overall measures, the model provides a close match to the adult subject data for both the time taken to complete the Tower and the number of constructions made in completing the Tower, as shown in Table 2. The model actually provides a close match to the adult subject data on seven out of nine measures of behavior (Jones, 1998), but only the two primary measures are reported here.

Insert Table 2 About Here

Performance on the Tower task is expected to show learning because each layer has the same characteristics, so problem solving on subsequent layers should be faster. The effect of learning on task performance can be examined from the time taken and the number of constructions made on a layer-by-layer basis. Figure 3 shows the mean time taken for the model (N=10) and subjects (N=5) to produce each layer, and the mean number of constructions made in producing each layer. There is a good correlation between the model and the adult subjects for the mean time to produce each layer ($r=0.96$; RMS error=4.1%), but not for the mean number of constructions made in producing each layer ($r=0.40$; RMS error=5.7%) because the curve for adults is flat. The RMS error is good for both measures.

One interesting feature of the model's times and construction attempts in Figure 3 is that they do not reduce after the size4 layer. For the adults, the measures are at their lowest for the size3 layer. This suggests that the model is performing at its optimum level after it completes the size4 layer, whereas adults are performing at their optimum level after completing the size3 layer.

Insert Figure 3 About Here

Modifying the Adult Model

The adult model, because it provides a good match to the adult behavior, can be modified to incorporate proposed theoretical mechanisms of development. The ensuing behavior can be compared to the behavior of seven year old children on the Tower task to test theoretical mechanisms of development.

Three independent modifications are made to the adult model. Two modifications implement capacity mechanisms and one implements strategy choice. The results of each modification are compared against the behavior of seven year old children on the Tower task taken from a recent study by Reichgelt, Shadbolt, Paskiewicz, Wood and Wood (1993). Children were tutored in completing the Tower task, and then given the Tower to build unassisted. The post-test performance of contingently tutored (Wood & Middleton, 1975) children (N=5) is re-analyzed here. That is, the children's data reported here refer to the second time the children completed the Tower, when they were unaided. Using tutored children is not considered problematic because the study from which the data is taken indicates that contingent tutoring is not needed from the age of six and upwards (Reichgelt et al., 1993). Comparisons between the modified model and the seven year old children will be the same as the comparisons that were made between the original model and adults.

Limited capacity models

There are several capacity based theories of development. Each theory explains capacity in a different way. For Pascual-Leone (1969), capacity is the number of knowledge structures that can be active at any one time. For Case (1985), capacity remains invariant with age, but through experience it can be used more efficiently (e.g., chunking knowledge means the same knowledge can be represented using less capacity). This allows more items to be processed simultaneously. For Halford (1993), capacity is measured in terms of number of dimensions, which increase with age. At any given stage, only knowledge structures of a specific complexity can be represented.

There are two natural ways to limit capacity in the adult model based on its cognitive architecture. First, the activation of blocks and block features must be above a retrieval threshold. Increasing retrieval threshold restricts the number of blocks and block features in working memory.

This corresponds to Pascual-Leone's capacity theory that a limited number of elements can be active at any one time. Second, limiting the number of conditions that a rule can have in its condition may mean that some strategies cannot be used, or have to be implemented as more than one rule (if possible). This corresponds to Case's and Halford's theories that there is a limit to the number of elements that can be processed simultaneously.

Increasing retrieval threshold.

The capacity mechanism described by Pascual-Leone was operationalised by increasing the retrieval threshold parameter in the adult model. In ACT-R, memory elements need to be above the retrieval threshold in order to be considered active, so increasing the parameter will reduce the number of active memory elements. This will mean less memory elements can be considered for matching in rules. In the adult model, the retrieval threshold was set to the ACT-R default of 0 (the parameter usually ranges from 0 to 3 for adults, C. Lebiere, personal communication, February 4, 1999). The retrieval threshold parameter was varied from 1 to 9 in units of 1 (the initial level of activation of items in memory is 10.0, but items are subject to decay on each cycle). The model provides the best match to the children at a retrieval threshold level of 6 (henceforth the RT6 model).

Table 2 shows that the RT6 model modifies the adult model to be more child-like in terms of the time taken to complete the Tower and the number of constructions made in completing the Tower. However, the RT6 model does not correlate well with the children for these measures when they are examined on a layer-by-layer basis. The correlation between the model and the children for the mean time to produce each layer is $r=0.51$ (RMS error=27.6%), and for the mean number of constructions made in producing each layer is $r=0.53$ (RMS error=26.8%).

The retrieval threshold modification has shown that limiting the number of memory elements that can be active at any one time can shift the behavior of the adult model towards that of seven year old children. However, the layer-by-layer data show that the RT6 model fails to match the layer-by-layer scores of the children.

Modifying the number of memory elements in the condition of a rule.

The developmental theories described by Case and Halford both specify a limit on the number of elements that can be processed simultaneously. In the adult model, the number of elements that can be processed simultaneously corresponds to the number of memory elements specified in the condition of a rule. Limiting this number represents a method of operationalising the capacity theories of Case and Halford.

The number of memory elements in the conditions of rules varies from 1 to 12 in the adult model. There are three methods by which the number of memory elements can be reduced: deleting the rule; splitting the rule into two or more sequential rules; chunking the memory elements such that the same information can be represented using less elements. The first two will be used when examining the effect of limiting the number of memory elements in the condition of a rule (the model does not implement the third as yet).

The limit of memory elements in the condition of a rule was varied from seven to ten. To implement the highest level (ten), three rules were deleted and a further four were split in two. To implement the lowest level (seven), three rules were deleted (the same ones as for the ten condition limit) and a further sixteen were split in two. Rules were only deleted when they could not be broken down into smaller rules. The model provides the best match to the children when the limit on conditions in rules is nine (henceforth the 9Cond model).

Table 2 shows that the 9Cond modification has enabled the model to closely resemble the children for both time taken to complete the Tower and number of constructions made in completing the Tower. However, when examining these measures on a layer-by-layer basis, the 9Cond model does not correlate well with the children for the mean time to produce each layer ($r=0.37$; RMS error=34.7%) or the mean number of constructions made in producing each layer ($r=0.44$; RMS error=31.2%). As with the RT6 model, the overall measures match the children well, but the layer-by-layer measures do not.

Strategy choice model

Siegler and Shipley (1995) suggest that children's strategy use changes with experience. This implies that older children are more able to select appropriate strategies than their younger counterparts. Strategy choice in children's addition has been examined in models before (Shrager & Siegler, 1998; Siegler & Shrager, 1984). Examining strategy choice here therefore represents a test of its applicability across domains.

The adult model implements strategies as rules, or sequences of rules. When more than one rule can be applied, the one with the highest strength is selected. The model subjects the strength of rules to an expected gain noise of 0.04, which means that on some occasions, rules which do not have the highest strength are selected. The noise is placed on every rule in the model, thereby influencing strategy choice for all of the various strategies involved in constructing the Tower. Increasing the value of the noise reduces the model's knowledge about which rules are the most effective. This enables the examination of the model's behavior when strategy choice is not as effective as it was for the adult model.

The expected gain noise that is placed on the strength of rules was increased in steps of one unit, starting at one. The range of values for the parameter normally lie between 0 and 1 for adults meaning the maximum setting can be open-ended for children. The model provides the best match to the seven year old children at an expected gain noise level of 6 (henceforth the EGN6 model).

Table 2 shows that the expected gain noise modification has enabled the behavior of the model to move closer to that of children for both time taken and constructions made. Figure 4 shows comparisons between the EGN6 model and the children for the mean time to produce each layer, and the mean number of constructions made in producing each layer. There is a good correlation between the model and the children for both measures ($r=0.99$ and $\text{RMS error}=3.0\%$ for time taken to construct each layer; $r=0.73$ and $\text{RMS error}=26.8\%$ for constructions made per layer). The modification enables a good fit on layer-by-layer measures.

An interesting characteristic that the EGN6 modification has provided is the increase in time to produce the size3 and size2 layers, which matches the children's behavior. By the time the EGN6 model begins the size3 layer, some aspects of the size3 and size2 blocks are not known (during the

task, decay has reduced their activation to below the retrieval threshold). This increases the likelihood of an incorrect strategy being selected because there is less opportunity for the conditions to be met of rules which represent good strategies. This is not as problematic for the adult model because it takes less time in completing the first three layers.

Insert Figure 4 About Here

The overall fit between the EGN6 model and the seven year old children is very good. There are no reliable differences between the EGN6 model and the children on any of the individual layer measures, for both timings and construction attempts. Modifying strategy choice has provided a good fit to the seven year old data.

Summary

Independent, theoretically motivated modifications to a computational model of adult performance (implemented within a cognitive architecture) can enable the model to approximate the behavior of seven year old children on the Tower task. Each modification represented an interpretation of a theoretical mechanism of development, changing the behavior of the model in different ways. Some modifications matched the data from seven year old's better than others. Changes to strategy choice provided the best fit to the children's data.

There is widespread belief that furthering our understanding of cognitive change is of critical importance to progressing developmental theory (Siegler, 1989; Sternberg, 1984). By investigating the role of different proposed mechanisms of development, the influence that each developmental mechanism has on behavior has been illuminated (albeit in a single domain).

The model and task used have also shown that subject data can be compared on multiple measures within in a developmental model, including timing data. Previous computational models of development were compared with subject data on one or two measures (e.g., McClelland & Jenkins, 1991), and rarely included timing data. Comparing model and subject data on a variety of measures

provides a more detailed match to the subject data, provides more faith that the processes modeled were correct, and suggests further ways the model can be improved.

With additional work, models of other ages and other tasks can be created. These models will be able to explain individual differences within age groups as well as to explain the progression between ages (in terms of differences between the models rather than transition mechanisms, for the moment). In both cases, the models should be able to highlight the knowledge differences or architectural changes that lead to the differences in behavior. The mechanisms which are shown to influence the behavior of the model then become the focus for any transition mechanism, because transition mechanisms must account for them. In this way, the degrees of freedom for transition mechanisms can be reduced (Simon & Halford, 1995).

The next step is to carry out further independent modifications to the model, and then interact the modifications by including two or more within the same model. The effects of interacting modifications are expected to be stronger than individual modifications, enabling the model to match the seven year old children's behavior better than the individual modifications. This type of work will help to more directly answer the question "What develops?".

References

- Anderson, J. R. (1993). Rules of the mind. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Baxter, G. D., & Ritter, F. E. (1996). Designing abstract visual perceptual and motor action capabilities for use by cognitive models. (Technical Report 36). CREDIT, Psychology Department: University of Nottingham.
- Case, R. (1985). Intellectual development: Birth to adulthood. Orlando, FL: Academic Press.
- Elman, J. L., Bates, E. A., Johnson, M. H., Karmiloff-Smith, A., Parisi, D., & Plunkett, K. (1996). Rethinking innateness: A connectionist perspective on development. Cambridge, MA: MIT Press.
- Gentner, D., Rattermann, M. J., Markman, A., & Kotovsky, L. (1995). Two forces in the development of relational similarity. In T. J. Simon & G. S. Halford (Eds.), Developing cognitive competence: New approaches to process modeling, 263-313. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Halford, G. S. (1993). Children's understanding: The development of mental models. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Jones, G. (1998). Testing mechanisms of development within a computational framework. Unpublished doctoral dissertation, University of Nottingham, Nottingham.
- Jones, G., & Ritter, F. E. (1997). Modelling transitions in children's development by starting with adults. In Proceedings of the European Conference on Cognitive Science (ECCS '97), 62-67. Manchester, UK: AISB.
- Jones, G., & Ritter, F. E. (1998). Initial explorations of modifying architectures to simulate cognitive and perceptual development. In Proceedings of the Second European Conference on Cognitive Modelling, 44-51. Nottingham University Press.
- Kail, R. (1991). Development of processing speed in childhood and adolescence. In H. W. Reese (Ed.), Advances in child development and behavior, (Vol. 23), 151-185. London: Academic Press.

Klahr, D., & Siegler, R. S. (1978). The representation of children's knowledge. In H. W. Reese & L. P. Lipsett (Eds.), Advances in child development and behavior, (Vol. 12). New York: Academic Press.

Klahr, D., & Wallace, J. G. (1976). Cognitive development: An information processing view. Hillsdale, NJ: Lawrence Erlbaum Associates.

McClelland, J. L., & Jenkins, E. (1991). Nature, nurture and connections: Implications of connectionist models for cognitive development. In K. VanLehn (Ed.), Architectures for intelligence. Hillsdale, NJ: Lawrence Erlbaum Associates.

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.

Murphy, C. M., & Wood, D. J. (1982). Learning through media: A comparison of 4-8 year old children's responses to filmed and pictorial instruction. International Journal of Behavioural Development, 5, 195-216.

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

Pascual-Leone, J. (1969). Cognitive development and cognitive style. Unpublished doctoral dissertation, University of Geneva.

Piaget, J. (1950). The psychology of intelligence. London: Routledge & Kegan Paul.

Reichgelt, H., Shadbolt, N. R., Paskiewicz, T., Wood, D. J., & Wood, H. (1993). EXPLAIN: On implementing more effective tutoring systems. In A. Sloman, D. Hogg, G. Humphreys, D. Partridge, & A. Ramsay (Eds.), Prospects for Artificial Intelligence, 239-249. Amsterdam: IOS Press.

Ritter, F. E., Baxter, G. D., Jones, G., & Young, R. M. (1999). Exploiting user interface management systems in cognitive modeling. Manuscript submitted for publication.

Shrager, J., & Siegler, R. S. (1998). SCADS: A model of children's strategy choices and strategy discoveries. Psychological Science, 9, 405-410.

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, 205-261. Hillsdale, NJ: Lawrence Erlbaum Associates.

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

Siegler, R. S. (1995). Children's thinking: How does change occur? In W. Schneider & F. Weinert (Eds.), Memory, performance and competencies: Issues in growth and development, 405-430. Hillsdale, NJ: Lawrence Erlbaum Associates.

Siegler, R. S., & Jenkins, E. A. (1989). How children discover new strategies. Hillsdale, NJ: Lawrence Erlbaum Associates.

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

Siegler, R. S., & Shrager, J. (1984). Strategy choices in addition and subtraction: How do children know what to do? In C. Sophian (Ed.), Origins of cognitive skills, 229-293. Hillsdale, NJ: Lawrence Erlbaum Associates.

Simon, T. J., & Halford, G. S. (1995). Computational models and cognitive change. In T. J. Simon & G. S. Halford (Eds.), Developing cognitive competence: New approaches to process modeling, 1-30. Hillsdale, NJ: Lawrence Erlbaum Associates.

Sternberg, R. J. (1984). Mechanisms of cognitive development: A componential approach. In R. J. Sternberg (Ed.), Mechanisms of cognitive development, 163-186. Prospect Heights, IL: Waveland Press.

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

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.

Young, R. M. (1976). Seriation by children: An artificial intelligence analysis of a Piagetian task. Basel: Birkhauser.

Author Note

The authors would like to thank Peter Cheng and Robert Siegler for useful comments during the preparation of this paper, Peter Cheng and Heather Wood for creating the initial version of Figure 1, and Hans Reichgelt, Nigel Shadbolt, Theresa Paskiewicz, and Heather Wood for collecting the videotaped data of seven year old children completing the Tower task.

Table 1

An example rule from the ACT-R Tower model.

```

(P grab-constrn-left-hand-not-over-constrn
  ;; IF the goal is to grab a construction using the left hand
  =goal>
    ISA                grab-constrn
    constrn            =constrn
    hand               left
    goal-done          nil
  ;; AND the left hand is NOT currently near the construction
  ;; Note: "-" represents NOT
  =left-hand>
    ISA                left-hand
    - is-over          =constrn
  ;; AND the model's eye is looking at the construction so that
  ;; it knows where to move the left hand to
  =focus>
    ISA                focus
    object              =constrn
  ==>
  ;; THEN create a goal to move the left hand to the construction
  ;; and push the goal onto the goal stack
  =newgoal>
    ISA                move-hand
    constrn            =constrn
    hand               left
  !push!              =newgoal
)

```

Table 2

Time taken and constructions made in completing the Tower, for the adults (N=5), the models (N=10), and the seven year old children (N=5). Standard deviations are shown in parentheses.

	Adults	Adult model	RT6 model	9Cond model	EGN6 model	Children
Time taken	126.6 s (34.0)	129.0 s (31.5)	160.2 s (35.6)	188.2 s (41.5)	166.3 s (37.9)	174.0 s (32.1)
Constructions made	22.8 (2.9)	23.1 (2.4)	25.8 (3.7)	29.6 (4.9)	25.1 (2.9)	27.6 (3.5)


Figure Captions

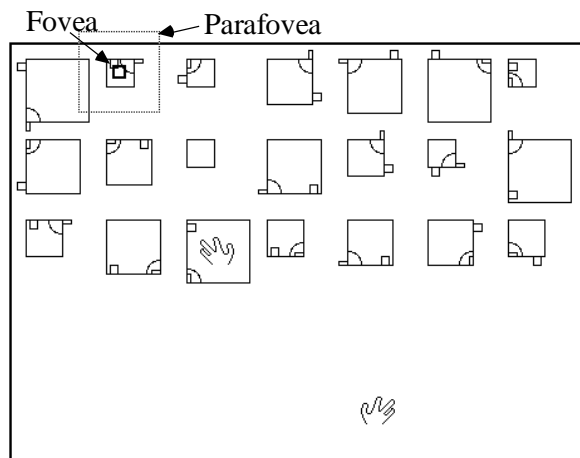
Figure 1. The blocks, left, that make up each layer, which are then stacked to create a tower, right.

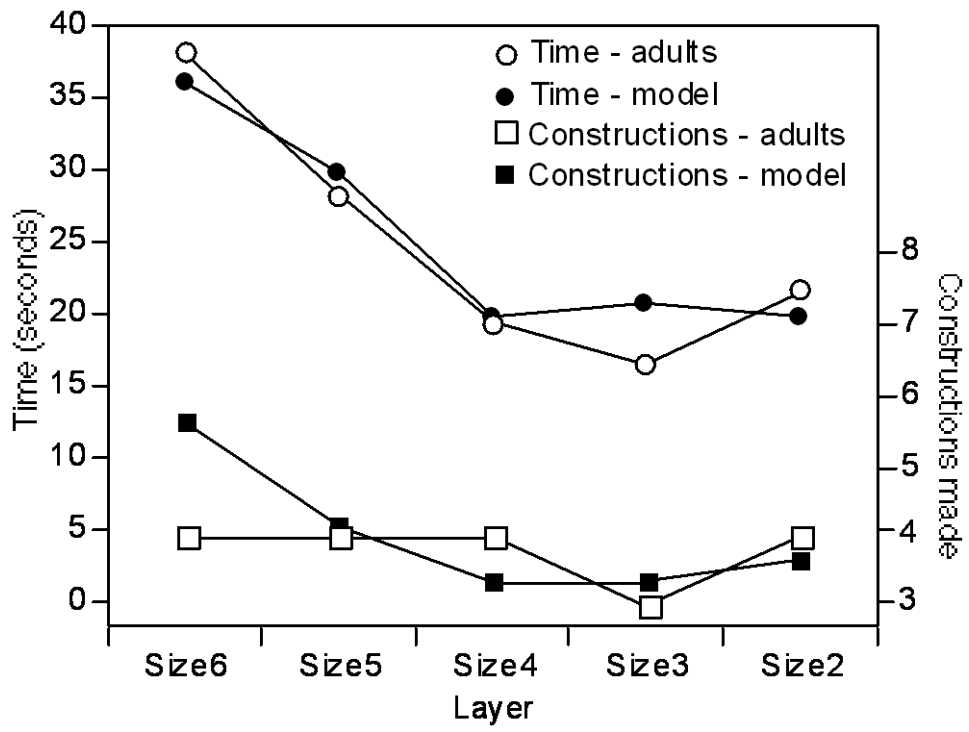
Figure 2. Graphical simulation of the Tower task.


Figure 3. Time taken and constructions made in completing each of the five layers, for the adults and the model (size6=largest layer blocks, size2=smallest layer blocks).

Figure 4. Time taken and constructions made in completing each of the five layers, for the seven year old children and the EGN6 model (size6=largest layer blocks, size2=smallest layer blocks).









¹ These include changes in thinking and/or representation on which the Piagetian concept of stage is based.