Inhibition and Young Children's Performance on the Tower of London Task

Frank D. Baughman (f.baughman@psychology.bbk.ac.uk)

Richard P. Cooper (r.cooper@bbk.ac.uk)

School of Psychology, Birkbeck, University of London,
London WC1E 7HX, UK

Abstract

There are relatively few explicit accounts detailing hypothesized mechanisms underlying executive functions. To understand the processes involved in problem solving at different stages in cognitive ability, we argue a mechanistic account is necessary. This paper outlines two initial computational models that simulate separately the performance of 3-4 year old and 5-6 year old children on the Tower of London reasoning task. We seek to capture the emerging role of inhibition in the older group. The basic framework is derived from Norman and Shallice's (1986) theory of willed and automatic action and Fox and Das' domino model (2000). Two strategies and a simple perceptual bias are implemented and comparisons reveal a good fit for the data of the two groups.

Introduction

It has been argued that our ability to reason, plan and solve novel and complex problems sets us apart from other animals. What are the processes or mechanisms that underlie these abilities? How do they function? Does some unitary component control them, or is control distributed? These questions are central to the study of *executive functions* (EF).

A substantial amount of research has focused on the role of three mechanisms of EF: (1) *inhibition of prepotent responses* ('inhibition'), (2) *shifting of mental sets* ('shifting') and (3) *updating of working memory* ('updating'). However, the exact conceptualizations of these underlying mechanisms and their levels of involvement in cognition have been a source of debate. Miyake, Friedman, Emerson, Witzki, Howerter and Wager (2000) offer one platform on which to ground these mechanisms, but studies that explicitly test the roles these mechanisms assume are lacking.

Drawing on a selection of research, the goal of this paper is to provide one detailed account of the possible role of inhibition in the Tower of London – a commonly used test of problem solving. We base our account of inhibition within a framework derived from two sources: (1) Norman and Shallice's theory of willed and automatic action (1986) and (2) the Domino model (Fox & Das, 2000). This paper demonstrates how a combination of their major features can lead to a broader theoretical framework within which the process of problem solving can be explored, and hypothesized mechanisms can be assessed more completely.

A number of developmental accounts of performance on the ToL exist (see e.g., Blair, Zelazo, & Greenberg, 2005) that are broadly supportive of a complex and dynamic interaction of mechanisms underlying high-level cognitive processes. Additionally, converging evidence for separable and differential levels of involvement of EF mechanisms, as espoused by Miyake et al., (2000) has also been claimed in a number of fMRI studies (see e.g., Sylvester, Wager, Lacey, Hernandez, Nichols, Smith & Jonides, 2003).

Consistent with these views on the development of EF mechanisms we adopt a position that assumes an association between age and the increased ability to inhibit. However, examining how these factors interact to support the development of mechanisms is not the focus of this paper.

Executive functions and Cognition

A consensus exists on the general definition of EF as a set of "general purpose control mechanisms that modulate the operation of various cognitive subprocesses and thereby regulate the dynamics of cognition" (Miyake et al, 2000, p. 50). However, at a more detailed level disagreement stems from the role given over to the 'executive'. A fundamental difference in the way EF has been characterized relates to the split between those theorists who argue a view of the executive as a *unitary* component (controlling and regulating behavior) and those viewing the processes as *diverse* (with the executive merely carrying out instructions it receives).[1]

Experimental evidence for Executive Functions

A number of tests exist that load heavily on different cognitive functions. Miyake et al. (2000) considers twelve different EF tests, including Wisconsin Card Sorting task, Stroop and the Tower of Hanoi. The assumption that these tests measure specific cognitive components is derived from observed dissociations of performance on them.

Inhibition is typically considered demonstrable by superior performance on tasks where an automatic or

---

[1] Zelazo, Reznick and Frye (1997) illustrate this distinction as the difference between the chief executive officer of a company and the executor of a will.

dominant response should be suppressed. For instance, this may be the successful inhibition of the tendency to process the semantics rather than the actual color of word items on the Stroop test. Conversely, a lack in the ability to inhibit is implied by poorer performance.

Theorists arguing both sides of the unity vs. diversity debate have used these tests to explicate the role of EF (see e.g., Rogers & Monsell, 1995; and Miyake et al., 2000). However, progress has been hindered by the multiple and somewhat "arbitrary and post-hoc" interpretations of resulting data (Miyake et al 2000, p53; see also Bull, Espy & Senn, 2004).

With the objective of clarifying some of these issues, Miyake et al. (2000) detail a study in which they find support for separable mechanisms of EF. Consistent with many previous studies and of special relevance to this paper were their findings that inhibition was most strongly associated to performance on the ToH than on a range of other EF tasks. Miyake et al. offer one plausible interpretation of these findings, reasoning that in the ToH one is influenced by the tendency to move *towards* greater perceptual similarity rather than move away. This interpretation fits with numerous other studies in which moves that take the configuration of the current state away from the goal state are described as counter-intuitive, or undesirable whilst in fact they are necessary for task completion (Gilhooly, 2002). We examine this position further within the ToL.

In this task, subjects are presented with an apparatus that consists of three colored balls and a board with three pegs of different lengths. The length of each peg constrains the number of balls that can be placed on it to 1,2, and 3 balls, respectively. They are then presented with a picture of the goal state and are asked to move each ball, one at a time to match the goal using the picture as a reference.

Problem solving and the Tower of London

Derived from the Tower of Hanoi and originally introduced by Shallice (1982), the ToL has become a popular tool for measuring problem solving abilities of children and adults with neurological impairments. Both tasks have been held to load heavily on EF and, in the case of the ToL in particular, on inhibition (Bull et al., 2004; Miyake et al, 2000). Figure 1 below provides examples of a) a 3-move tower-ending and b) a 3-move flat-ending problem.[2]

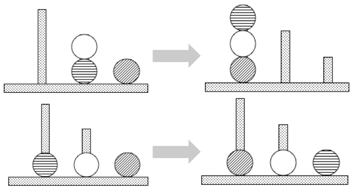

Figure 1. Tower-ending and flat-ending 3-move ToL problems

From the research on the ToL it is possible to extract two strategies that young children may use in problem solving. These form specific components within the computational models developed here. The strategies are (1) an *immediate-hit* strategy (the tendency to place a ball in its target position immediately if the target-position is free and the target-ball is free to move) and (2) a *one-move look-ahead* strategy (the ability to plan moves up to one-move away). See e.g. Bull et al., (2004) and Goel, Pullara and Grafman (2001).

Of some importance in developing these models is the need to specify the criteria that are to be used to evaluate their performance. The ToL is typically measured using data on (a) whether the correct solution was reached and (b) whether the solution was achieved within the minimum number of moves (and quite often this does not include a record of the actual number of moves made; see e.g., Bull, 2004). However, these may not be the best nor the only measurements of performance. Specifically, two features of young children's behavior appear to have often been over-looked. These are *rule breaks* (e.g., moving two balls at a time) and *partial-completion* (e.g., stopping when only one or two balls are in their correct positions). These measurements are incorporated into the current study to allow a more detailed understanding of young children's abilities in problem solving and provide a fuller model of behavior.

A framework of behavioral control (SAS and CS)

Norman and Shallice's (1986) theory of willed and automatic action is perhaps one of the best-known examples to embody the diversity view of cognition. Divisible into two distinct but significantly related processes that operate according to specific parameters, it comprises a contention scheduling (CS) system and a supervisory attentional system (SAS). Briefly, the CS organizes routine behaviors in the form of schemas and is characterized by low-level, predominantly autonomous processes that control everyday actions. The SAS imposes a heavy top-down influence by way of generating goals, creating schemas for CS to carry out and monitoring behavior. Problem solving and behavior in general is thus held to be the product of the influences of these two interrelated systems, with the SAS more involved in novel tasks but the CS taking over when tasks become familiar. The depth and breadth of behavior the CS-SAS theory is intended to account for makes it a suitable starting point for modeling the ToL and its key characteristics find important parallels within several approaches to problem solving (see, Glasspool, 2005).

The domino model

Fox and Das' (2000) domino model (see Figure 2) represents a highly organized system for decomposing elements of a problem. While it is not assumed to be a model of human cognition, it has been used extensively in AI work on expert systems (see Fox & Das, 2000)

---

[2] Though variation in performance on these problems may reveal important differences in the kinds of processes used on tasks consisting of clear vs. ambiguous sub-goals, they are not the focus of this paper (see Baughman, 2005).

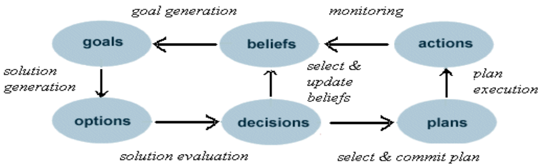

Figure 2. The generalized domino model of Fox & Das (2000)

The domino's relevance here is due to both general and specific similarities it shares with the SAS component of Norman and Shallice's (1986) model (for detailed review see Glasspool, 2005; Glasspool & Cooper, 2002). It thus offers a reasonable framework within which the SAS and CS processes of Norman and Shallice's theory can be fleshed out more completely.

Computational approaches to problem solving

Computational approaches using Soar (Newell, 1990) and ACT-R (e.g., Anderson & Lebiere, 1998) have contributed greatly to the study of a wide range of cognitive behavior. However, their use has also drawn criticisms on a number of theoretical and technical fronts (see e.g., Cooper & Shallice, 1995 and Altman & Trafton, 1999, respectively). In this paper, we seek to limit the number of underlying theoretical assumptions in modeling the ToL, and choose to adopt COGENT (Cooper & Fox, 1998; Cooper, 2002). In contrast to Soar and ACT-R, COGENT does not specify any particular theoretical architecture within which models must be placed and therefore allows only aspects of theory that are deemed specially relevant to be included. Furthermore, COGENT allows the development of symbolic and connectionist processes within the same model. These features make it an attractive environment within which to develop the mix of automatic low-level processes assumed within the CS and the higher-level processes within the SAS.

Modeling the Tower of London

We describe two related models of behavior on the ToL task. The models attempt to capture critical features of the behavior of two groups of children, a group of 3-4 year olds (the Younger Child model) and a group of 5-6 year olds (the Older Child model). The target behavioral data come from a study in which children's performance on two different types of ToL, tower-ending and flat-ending was compared (Waldau, 1999). As in the original behavioral study each model is run on six tasks for a total of 17 pseudo-subjects. The models take as their problem-set the 3 tower-ending and 3 flat-ending tasks used in the original study. A visual representation of one problem is given in figure 3. For present purposes comparisons of these data are collapsed across problem type to highlight general patterns of each models' performance.

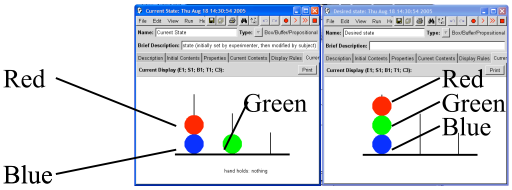
Current State: Thu Aug 18 14:30:54 2005

Name: Current State

Type: Box/Buffer/Propositional

Brief Description: state (initially set by experimenter, then modified by subject)

Current Display (E1; S1; B1; T1; C3):

hand holds: nothing

Desired state: Thu Aug 18 14:30:54 2005

Name: Desired state

Type: Box/Buffer/Propositional

Brief Description:

Current Display (E1; S1; B1; T1; C3):

Figure 3. Current and Desired states

The 'Younger Child' model

The framework adopted for each model and displayed in figures 4 and 7 shows the influence of the domino model and Norman and Shallice's theory at the level of the subject and within the Contention Scheduling process, respectively.

The Subject. Displayed in Figure 4, the Subject model derives representations of the Current and Desired states (see figure 3) through Perception of World. This process extracts simple properties of the task and maintains representations of the Current and Desired states in Working Memory.

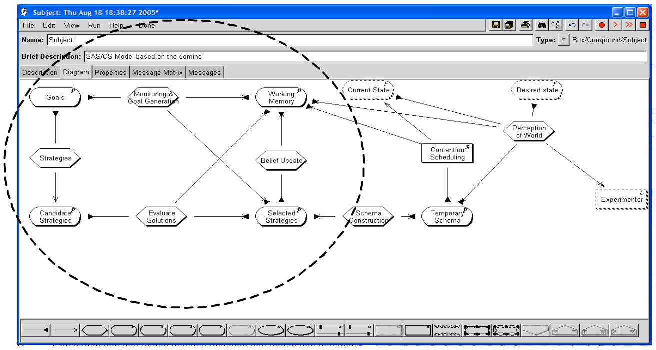
Subject: Thu Aug 18 18:38:27 2005*

Name: Subject

Type: Box/Compound/Subject

Brief Description: SAS/CS Model based on the domino

Figure 4. The processes within Subject

Problems are recognized as problems by **Monitoring and Goal Generation** (Monitoring), if the Current and Desired states do not exactly match. When there is a discrepancy between these two states, a message is produced and sent to **Goals** that triggers the use of existing strategies aimed at reducing the difference. Strategies delivers the immediate-hits and one-move look ahead strategies to **Candidate Strategies**, which are then analyzed by **Evaluate Solutions**.

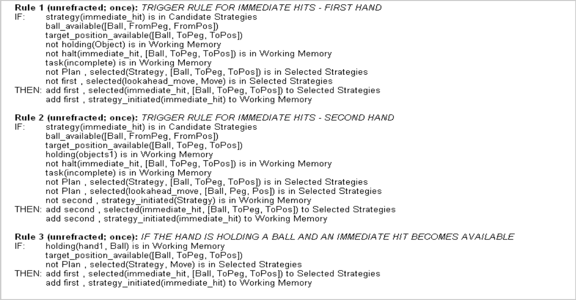
```
Rule 1 (unrefracted; once): TRIGGER RULE FOR IMMEDIATE HITS - FIRST HAND
IF:     strategy(immediate_hit) is in Candidate Strategies
        ball_available([Ball, FromPeg, FromPos])
        target_position_available([Ball, ToPeg, ToPos])
        not holding(Object) is in Working Memory
        not halt(immediate_hit, [Ball, ToPeg, ToPos]) is in Working Memory
        task(incomplete) is in Working Memory
        not Plan , selected(Strategy, [Ball, ToPeg, ToPos]) is in Selected Strategies
        not first , selected(lookahead_move, Move) is in Selected Strategies
THEN:   add first , selected(immediate_hit, [Ball, ToPeg, ToPos]) to Selected Strategies
        add first , strategy_initiated(immediate_hit) to Working Memory

Rule 2 (unrefracted; once): TRIGGER RULE FOR IMMEDIATE HITS - SECOND HAND
IF:     strategy(immediate_hit) is in Candidate Strategies
        ball_available([Ball, FromPeg, FromPos])
        target_position_available([Ball, ToPeg, ToPos])
        holding(objects1) is in Working Memory
        not halt(immediate_hit, [Ball, ToPeg, ToPos]) is in Working Memory
        task(incomplete) is in Working Memory
        not Plan , selected(Strategy, [Ball, ToPeg, ToPos]) is in Selected Strategies
        not Plan , selected(lookahead_move, [Ball, Peg, Pos]) is in Selected Strategies
        not second , strategy_initiated(Strategy) is in Working Memory
THEN:   add second , selected(immediate_hit, [Ball, ToPeg, ToPos]) to Selected Strategies
        add second , strategy_initiated(immediate_hit) to Working Memory

Rule 3 (unrefracted; once): IF THE HAND IS HOLDING A BALL AND AN IMMEDIATE HIT BECOMES AVAILABLE
IF:     holding(hand1, Ball) is in Working Memory
        target_position_available([Ball, ToPeg, ToPos])
        not Plan , selected(Strategy, Move) is in Selected Strategies
THEN:   add first , selected(immediate_hit, [Ball, ToPeg, ToPos]) to Selected Strategies
        add first , strategy_initiated(immediate_hit) to Working Memory
```

Figure 5. Excerpt of rules pertaining to the Immediate-hit strategy

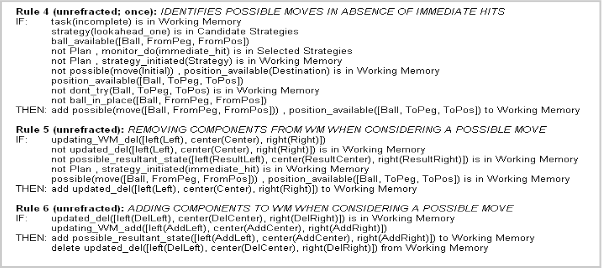

```
Rule 4 (unrefracted; once): IDENTIFIES POSSIBLE MOVES IN ABSENCE OF IMMEDIATE HITS
IF:     task(incomplete) is in Working Memory
        strategy(lookahead_one) is in Candidate Strategies
        ball_available([Ball, FromPeg, FromPos])
        not Plan , monitor_do(immediate_hit) is in Selected Strategies
        not Plan , strategy_initiated(Strategy) is in Working Memory
        not possible(move(Initial)) , position_available(Destination) is in Working Memory
        position_available([Ball, ToPeg, ToPos])
        not dont_try(Ball, ToPeg, ToPos) is in Working Memory
        not ball_in_place([Ball, FromPeg, FromPos])
THEN:   add possible(move([Ball, FromPeg, FromPos])) , position_available([Ball, ToPeg, ToPos]) to Working Memory

Rule 5 (unrefracted): REMOVING COMPONENTS FROM WM WHEN CONSIDERING A POSSIBLE MOVE
IF:     updating_WM_del([left(Left), center(Center), right(Right)])
        not updated_del([left(Left), center(Center), right(Right)]) is in Working Memory
        not possible_resultant_state([left(ResultLeft), center(ResultCenter), right(ResultRight)]) is in Working Memory
        not Plan , strategy_initiated(immediate_hit) is in Working Memory
        possible(move([Ball, FromPeg, FromPos])) , position_available([Ball, ToPeg, ToPos]) is in Working Memory
THEN:   add updated_del([left(Left), center(Center), right(Right)]) to Working Memory

Rule 6 (unrefracted): ADDING COMPONENTS TO WM WHEN CONSIDERING A POSSIBLE MOVE
IF:     updated_del([left(DelLeft), center(DelCenter), right(DelRight)]) is in Working Memory
        updating_WM_add([left(AddLeft), center(AddCenter), right(AddRight)])
THEN:   add possible_resultant_state([left(AddLeft), center(AddCenter), right(AddRight)]) to Working Memory
        delete updated_del([left(DelLeft), center(DelCenter), right(DelRight)]) from Working Memory
```

Figure 6. Excerpt of rules pertaining to the one-move look-ahead strategy

Evaluate Solutions

Evaluate Solutions provides intensive processing of information represented within Working Memory and to a lesser extent Selected Strategies. The primary objective for this process is to evaluate the outcome of proposed solutions, or moves.

Evaluate Solutions is responsible for identifying immediate-hits (see Figure 5) and look ahead moves (see Figure 6). In the event where none exist, Evaluate Solutions starts a process whereby possible moves are proposed to Working Memory. Evaluate Solutions calculates what the *resultant state* would be if that possible move was actioned. If an immediate-hit is possible given a resultant state, the possible move is initiated.[3] If a possible move does not yield an outcome whereby an immediate hit is possible then it is temporarily 'black listed' in working memory and another possible move is explored. Once the 'decision' has been made to move a ball, automatic processes within CS take over.

The left-hand side of the model (highlighted by the dashed circle, Figure 4) is thus given over to 'decision-making' as strategies are proposed and evaluated. For each Selected Strategy action schemas are created (via **Schema Construction**) and fed into **Temporary Schemas**, serving to excite elements within Contention Scheduling.

Contention Scheduling

Figure 7 illustrates the processes within Contention Scheduling in which action schemas (such as 'pick up green' and 'put down on centre peg') are ultimately produced. Automatic processes within **Trigger Schemas** operate on (1) balls that can be moved and (2) pegs that have space, by reading from the Current State.

In the absence of strategies and hence schemas, processes within CS also organizes problem-elements according to the *configural* similarity between the Current and Desired states. Configural similarity refers to the degree to which the Current and Desired States appear similar. This bias to move a ball to a location that matches the overall configuration of the goal state, but not necessarily the correct color position mirrors findings from experimental data and serves to allow basic responses to take place in more complex situations where existing strategies do not appear suited.

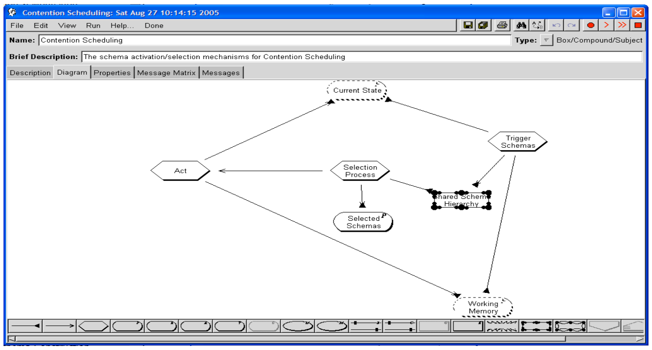

Name: Contention Scheduling
Type: Box/Compound/Subject
Brief Description: The schema activation/selection mechanisms for Contention Scheduling

Figure 7. The contention scheduling process

For each ball that can be moved, an excitatory message is sent to the corresponding node in **Shared Schema Hierarchy** (figure 8). When the level of activation for a node reaches its threshold (0.75) that action is selected and carried out through **Act**. The contents of Current State and Working Memory are then updated to reflect the new positions of balls and the process of determining the next possible move begins.

In addition to maintaining representations of the environment in Working Memory, Perception of World functions to provide information that allows the level of configural similarity to be determined. In this way a number of possible action schemas can emerge that are not based upon higher-level strategies, but on direct perceptual features.

For each ball that is out of place, a schema is proposed to move it. E.g., in Figure 3 above, the blue and red balls are in correct configural positions (only blue in its actual color position) but the green ball is out of place. Thus in this example, only one schema would be generated – to move the green to the left peg.

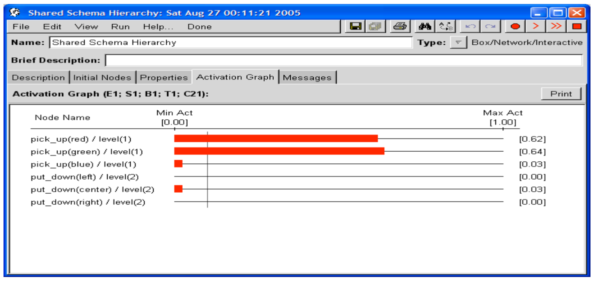

Name: Shared Schema Hierarchy
Type: Box/Network/Interactive
Brief Description:

Activation Graph (E1; S1; B1; T1; C21):

Node Name	Min Act [0.00] – Max Act [1.00]
pick_up(red) / level(1)	[0.62]
pick_up(green) / level(1)	[0.64]
pick_up(blue) / level(1)	[0.03]
put_down(left) / level(2)	[0.00]
put_down(center) / level(2)	[0.03]
put_down(right) / level(2)	[0.00]

Figure 8. The interactive-activation network of CS

A linear increase of 1:0.22 was used to establish the activation values associated with each level of configural similarity. For example, based on Fig 3 above, the schema to move the green to the left peg would produce an excitation of 0.66 within Shared Schema Hierarchy.

---

[3] This tendency to abandon full searches of the problem space for other beneficial moves is consistent with the literature (Gilhooly, 2002).

Model summary

The performance of this model is influenced by the existence of two strategies and a simple perceptual bias to match the current configuration to the Goal State. The fit with the experimental data is good for the 3-4 year olds on all criterion measurements (see Table 1).

Table 1. Summary of behavior of older children and model 1

	3-4 yr olds	Model 1
Config (%)	95.83	100.00
Colors (%)	66.32	65.74
Rule breaks (%)	52.08	44.11
Avg no. moves	11.26	6.8

Two interrelated processes may account for the simulated number rule-breaks. Firstly, in the absence of strong lateral inhibition between nodes, activation of two nodes e.g., ‘pick-up red’ and ‘pick-up blue’ can reach threshold and result in both actions being taken. As a ball is picked up, it may reveal another ball to which a second strategy applies. Hence, the schema to pickup the second ball can reach threshold before the first ball has been placed.

Secondly, partial completion (or, the mixture of lower number of correct colors and high number of correct configurations) is explained as the result of a combination of effects of the immediate-hit strategy and the direct, configural bias within the CS. Both of these influences are concerned with immediate perceptual properties of the Current State. In the case of the former, the strategy is concerned purely with placing a ball in its target position and does not process the placement of other balls. In the latter, only configural properties are processed. In the course of problem solving the perceptual features of a problem are present *before* the results of processing of the various strategies have been carried out (i.e., these more intensive processes take longer to return proposed moves). The CS (containing the bias for configural similarity) is dependent on perceptual information only and so has an early advantage at influencing the selection of schemas. Thus, the direct the influence of CS goes unchecked.

The lack of co-ordination between CS and supervisory processes suggests the need for greater monitoring and control to *inhibit* the influence of simple and direct perceptual biases that originate from the CS. In the next model we introduce a mechanism of inhibition to simulate the performance of the 5-6year olds.

The ‘Older Child’ model

The second model extends the first by interrupting the combined influence of immediate-hits and configural bias, thereby enabling a greater degree of influence from higher-level strategies.

The role of Inhibition

In a strong top-down manner, Monitoring and Goal Generation intervenes directly in the process of problem solving by temporarily suppressing a move if the position to which the ball is to be placed is *above* another ball. Moves proposed by the immediate-hit strategy (via Evaluate Solutions) and moves influenced by the configural bias are halted until further checks are carried out.

This difference between the Younger Child model and Older Child model is embodied as a single rule within Monitoring and Goal Generation (see Figure 9 below). If its condition is met a more detailed examination of the positions of other balls in the current state is triggered. If this reveals the ball under the target position for the immediate-hit is not in place, the strategy is terminated and a new move considered.

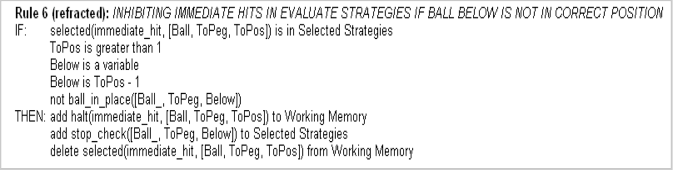

Rule 6 (refracted): *INHIBITING IMMEDIATE HITS IN EVALUATE STRATEGIES IF BALL BELOW IS NOT IN CORRECT POSITION*

```
IF:     selected(immediate_hit, [Ball, ToPeg, ToPos]) is in Selected Strategies
        ToPos is greater than 1
        Below is a variable
        Below is ToPos - 1
        not ball_in_place([Ball_, ToPeg, Below])
THEN:   add halt(immediate_hit, [Ball, ToPeg, ToPos]) to Working Memory
        add stop_check([Ball_, ToPeg, Below]) to Selected Strategies
        delete selected(immediate_hit, [Ball, ToPeg, ToPos]) from Working Memory
```

Figure 9. Inhibiting the immediate-hit strategy if the ball under is not in place

Model summary

The Older Child model adds an important feature that serves to inhibit actions based on simple and direct perceptual biases. These biases are suppressed via a rule that triggers a deeper search of the problem state. The effects of this mechanism within the Monitoring process is in reducing the chances of the model being ‘led astray’ by superficial characteristics of the problem and increasing the proportion of balls being placed in their correct color position.

Table 2. Summary of behavior of older children and model 2

	5-6 yr olds	Model 2
Config (%)	96.05	100.00
Colors (%)	94.74	95.05
Rule breaks (%)	21.92	30.47
Avg no. moves	9.79	7.17

A comparison of the behavior of this model and older children is given in Table 2. Overall, it appear that inhibition holds considerable weight on the overall behavior of the Older Child model. We argue that while it is an essential component of older children’s performance, the final outcome on each task is reliant on a number of other processes. This interpretations fit well both with diversity accounts and studies emphasizing a strong involvement of inhibition on tasks of EF, including the ToL (Miyake et al, 2000).

General discussion

The models presented here integrate a number of aspects from the work by Norman and Shallice’s and Fox and Das to achieve a framework capable of testing a candidate role of inhibition on the ToL. Overall, these models demonstrate a good fit of the data for 3-4 and 5-6 year olds.

The effect of the simple perceptual bias enables moves to be made towards the configuration of a goal. Determining which ball to move and which peg to place it at, is a result of competition within an interactive-activation based network; with the level of configural similarity governing the amount of excitation competing nodes receive. If configural similarity is high and either no strategies apply, or more extensive processing is required to apply a strategy, this bias has a stronger possibility of influencing a move to configuration, increasing the chances of only a partially complete solution.

In the Older Child model one view of the possible role that inhibition may play in problem solving on the ToL was established. This view was built on conceptualizations offered by Miyake et al (2000). The implementation given to inhibition within the model described here accounts for a shift in performance of the model, from one resembling the behavior of 3-4 year olds to one resembling the behavior of 5-6 year olds.

Furthermore, rather than operating under a single mechanism the behavior of these models is determined by a range of processes. Although in the second model, the role given to inhibition is instrumental in accounting for specific differences between the Younger Child and Older Child models, influences of both strategies and the perceptual bias converge to affect performance. Thus, these models strongly favor diversity views of EF.

The work presented here is consistent with the view that younger children's poorer performance on the ToL is a product of their failure to inhibit simpler strategies. In contrast to the view that younger and older children possess qualitatively different cognitive strategies these models demonstrate that a lack of ability to inhibit may mask the existence of more complex skills.

The account offered here is a functional one. In this paper we have demonstrated that a computational implementation of inhibition can explain differences in performance between younger and older children. However, it remains to be explained what the mechanisms are that drive the development of this and other mechanisms that underlie cognitive development.

References

Anderson, J. R., & Lebiere, C. (1998*). Atomic Components of Thought.* Maywah, NJ: Erlbaum.

Altmann, E. M. & Trafton, J. G. (1999). Memory for Goals: An Architectural Perspective. *Proceedings of the twenty first annual meeting of the Cognitive Science Society,* 19–24. Hillsdale, NJ: Erlbaum.

Baughman, F.D. (2005). The Role of Inhibition in Young Children's Performance on the Tower of London: A Computational Study. *Unpublished Masters thesis,* University of London.

Blair, C., Zelazo, P.D., & Greenberg. M.T. (2005). The Measurement of Executive Function in Early Childhood. *Developmental Neuropsychology, 28(2)*, 561–571.

Bull, R., Espy, K. A., & Senn, T. E. (2004). A Comparison of Performance on the Towers of London and Hanoi in Young Children. *Journal of Child Psychology and Psychiatry, 45:4*, 743–754.

Cooper, R.P. & Fox, J. (1998). COGENT: A visual design environment for cognitive modelling. *Behaviour Research Methods, Instruments, & Computers, 30*, 553–564.

Cooper, R.P. and Shallice, T. (1995). Soar and the case for unified Theories of Cognition. *Cognition, 55,* 115–149.

Cooper, R.P. (2002). *Modelling High-Level Cognitive Processes*. Lawrence Erlbaum Associates, Mahwah, NJ.

Fox, J. & Das, S. (2000). *Safe and Sound: Artificial Intelligence in Hazardous Applications.* MIT press, Cambridge, MA.

Gilhooly, K. J. (2002). *Thinking: Directed, Undirected and Creative*. Academic Press.

Glasspool D.W. (2005). The integration and control of behaviour: Insights from neuroscience and AI. In D. N. Davis (Ed.) *Visions of Mind: Architectures for Cognition and Affect* (pp. 208–234). Hershey, PA: IDEA Group.

Glasspool, D.W. & Cooper, R.P. (2002). Executive Processes. In Cooper, R.P. (ed.) *Modelling High-Level Cognitive Processes* (pp. 313–362). Lawrence Erlbaum Associates, Mahwah, NJ.

Goel, V., Pullara, S. D., & Grafman, J. (2001). A computational model of frontal lobe dysfunction: Working memory and the Tower of Hanoi task. *Cognitive Sciences, 25,* 287–313.

Miyake, A., Friedman, N.P., Emerson, M.J., Witzki, A.H., Howerter, A., & Wager, T.D. (2000). The unity and diversity of executive functions and their contributions to complex frontal lobe tasks: A latent variable analysis. *Cognitive Psychology, 41,* 49–100.

Newell, A. (1990). *Unified Theories of Cognition*. Cambridge MA.: Harvard University Press.

Norman, D.A., & Shallice, T. (1986). Attention to action: willed and automatic control of behaviour. In R. Davidson, G. Schwartz, & D. Shapiro (Eds.), *Consciousness and self regulation* (Vol. 4). Plenum: New York.

Shallice, T. (1982). Specific impairments of planning. *Philosophical Transactions of the Royal Society London B 298*, 199–209.

Sylvester, C.Y.C., Wager, T. D., Lacey, S. C., Hernandez, L., Nichols, T. E., Smith, E. E., & Jonides, J. (2003). Switching attention and resolving interferences: fMRI measures of executive functions. *Neuropsychologia, 41*, 357–370.

Rogers, R. D., & Monsell, S. (1995). Costs of a predictable switch between simple cognitive tasks. *Journal of Experimental Psychology: General, 124*, 207–231.

Waldau, R. (1999). A developmental study of problem solving in children aged 3–6: Development of planning strategies. *Unpublished undergraduate thesis,* University of London.

Zelazo, P. D., Carter, A., Reznick, J. S., & Frye, D. (1997). Early development of executive function: A problem-solving framework. *Review of General Psychology, 1,* 198–226.