Developing Problem Solving Competence: A Distributed Model and a New Class of Strategies With the Tower of Hanoi Task

Bernadette Guimberteau (Bern0@mac.com)
Cognition and Development, School of Education
4533 Tolman Hall, University of California
Berkeley, CA 94720 USA

Abstract

The ability to focus on the largest disk of a pyramid at the outset and to define largest-disk subgoals constitute two essential aspects in most known strategies and models of problem solving with the Tower of Hanoi. Yet, those abilities are typically assumed by existing accounts. This paper presents a distributed model, which learns to focus on the largest disk of a pyramid and set subgoals to move largest disks. The model exhibits a capacity to solve 4- and 5-disk Tower of Hanoi versions optimally and to evolve toward more competent behavior. Moreover, the fit between this analysis and the data from Anzai & Simon (1979) is excellent. The present model provides a new interpretation of those data: the subject's learning is due to the acquisition of task-specific afffordances and of difference-reduction strategies that are affordance-driven. The above analysis defines a new class of Tower of Hanoi strategies – based on a problem solver's capacity to define and use task-specific affordances. The mechanisms proposed by the model can be used to examine the distributed nature of learning and problem solving in other tasks as well.

Introduction

This paper presents a distributed model of the development of problem solving strategies with the Tower of Hanoi puzzle. Problem solving strategies constitute a central theme in the study of adaptive cognitive behavior: However, cognitive science has been preoccupied mostly with their identification and with the elucidation of their role in problem solving (e. g., Altmann & Trafton, 2002; Karat, 1982; Simon, 1975), rather than with their origin.


Figure 1: The Tower of Hanoi Task: Initial State.

The Tower of Hanoi is a classical task in cognitive science. It begins with a state where a pyramid (or stack) of *k* disks is stacked on one peg (the source peg). The disks vary in size, as shown in Figure 1, which uses a five-disk version of the task with the disks standing on the source peg A. The pyramid needs to be transferred to another peg (the goal peg) by moving only one disk at a time and by placing only a smaller disk on top of another one.

When solving that puzzle, considering the largest disk of the pyramid at the outset constitutes a strategically useful approach because the movement of that disk is the most constrained: the task rule stating that a disk can go on top of another one only if it is smaller prevents bigger disks from being stacked on top of smaller ones, limiting their moving options. More generally, setting largest-disk subgoals - subgoals to move the successively largest disks of a pyramid to the goal peg (first the largest, then the next-largest, etc.) - offers a productive avenue to solve the puzzle, due to the constraints placed by the task rules on those disks.

The ability to consider the largest disk of a pyramid and its bigger disks represents a given for major strategies and problem solving models of the Tower of Hanoi task. For example, most known strategies set largest-disk subgoals. This is certainly true of the recursive subgoaling strategy (e.g., Altmann & Trafton, 2002; Anderson, Kushmerick & Lebiere, 1993; Simon, 1975). Given a pyramid of disks, that strategy assigns pegs beginning with the largest disk, and proceeds with the next largest until the smallest disk is moved. However, even models that are meant to be simpler – that is, not based on recursive properties – often require the ability to focus on the largest disk of a pyramid. For example, Simon (1975)'s perceptual strategies operate by identifying the largest disk not yet on the goal peg and by defining successive steps to get that disk to that peg. Karat (1982)'s problem solving model is also largest -disk driven – the model sets the subgoal to send to the goal peg the largest-disk not already on that peg, applying general search heuristics (e.g., avoiding to move the smallest disk twice in a row) until that goal is attained.

Even though most known strategies and models start out with a largest-disk emphasis, data do exist suggesting that the ability to focus on the largest disk of a pyramid may not always be present at the outset: it develops through interactions with the task instead. The problem solving protocol from Anzai & Simon (1979) provides a nice illustration. The participant in that study exhibited the capacity to focus on the largest disk of a five-disk pyramid toward the end of her first problem solving attempt only – not spontaneously at the beginning – and in a state where that disk was the only one on the source peg, not at the bottom of the stack. Moreover, she used largest-disk subgoals during her second problem solving attempt – not in the first one.

If the ability to focus on largest disks develops through experience with the task, then several questions arise: What are the mechanisms underlying this development? Do they shape the acquisition of a largest-disk focus exclusively or other aspects of problem solving as well during that learning process? What happens to those mechanisms once the ability to set largest-disk subgoals is acquired? Do they simply disappear, or do they still play a role in the definition of new strategies? Those are key issues for research seeking to understand strategy acquisition and the emergence of problem solving competence.

The above questions remain open. For example, analyses of Anzai & Simon (1979)'s problem solving protocol (Anzai & Simon, 1979; VanLehn, 1991) have not made clear the exact mechanisms behind the subject's ability to develop a focus on the largest disk. Moreover, the characterization of largest-disk subgoals in the second episode of that protocol has been described as occurring along with the Selective Search strategy identified in the first episode (Anzai & Simon, 1979), but this characterization has been questioned: Selective Search requires no subgoaling or planning and should produce brief comments, not verbal indicators showing that the subject is struggling (VanLehn, 1991).

The above observations suggest that models are needed to explicate the mechanisms behind the origin of a largest-disk focus, the setting of largest-disk subgoals, and the possible contribution of those mechanisms to the definition of problem solving strategies. Such a model is presented in what follows. It proposes that largest-disk focus and largest-disk subgoals develop from affordance-driven behavior, which also shapes other aspects of problem solving during that development. This hypothesis is suggested by the fact that inexperienced users tend to rely on affordances during task performance (Norman, 1988; 1993). A major purpose in building the present model is to explore the viability of the view that there exists optimal problem solving strategies that are affordance-driven. Given that this model seeks to account for largest-disk subgoaling – a capacity often associated with optimal problem solving strategies – it should be possible to define and observe optimal problem solving strategies structured around the affordance-driven mechanisms proposed here.

A Sketch of The Model

This section presents the main aspects of the model. They fall into three categories: affordances, difference-reduction strategies, and learning mechanisms.

Affordances

The concept of affordances refers here to perceived properties regarding the use of Tower of Hanoi objects (especially single disks and stacks of disks) – following Norman (1988)'s definition of affordances:

"…The perceived and actual properties of the thing, primarily those fundamental properties that determine just how the thing could possibly be used...." (Norman, 1988, p.9)

In the model, move affordances – the property of movability (perceived and actual) of Tower of Hanoi objects – constitute an essential concern for problem solvers because task rules restrict the move of certain objects (e.g., large disks) and facilitate that of others (e.g., the smallest disk).

Two classes of affordances are distinguished in the analysis: general and task-specific. The former refers to general properties that offer common, non task-specific operational clues, while the latter emerges in relation to task constraints, providing operational clues that integrate such constraints. An example of a general affordance is a move property attributed to disks: an unblocked disk affords moving. That affordance – called "Move: Movable_Disk" hereafter – is based on knowledge that is not task-specific: a small, unblocked object that can be grabbed affords moving. Three examples of task-specific affordance are that a stack of disks affords emptying under task rules ("Empty: Disk_Stack"), that the largest disk of a stack affords moving the least ("Largest_Disk" affordance), and that a small, movable disk affords moving more than a bigger, movable disk ("Relative_Move"). Those affordances emerge as the task constraint that requires moving only smaller disks on top of bigger ones is being followed.

The above classes of move affordances play an essential role in the model. One of their functions is to "externalize" a problem solver's internal, memorized task rules. For example, the Move: Movable_Disk affordance externalizes the internal rule to move only one disk at a time because the former embodies the latter: by following the property that an unblocked disk affords moving and by moving that disk, one also follows the task rule to only move one disk at a time. In other words, using the affordance applies the task rule. Another important function of affordances in the model is that they give rise to affordance-driven strategies that reduce differences toward the task goal – providing a measure of progress. For example, the Empty: Disk_Stack affordance supports the discovery of a strategy consisting in making a move that contributes to emptying the stack. Such a strategy provides a way to get closer to the task goal.

Difference-Reduction Strategies

The difference-reduction strategies considered in the model are elementary. They specify necessary but not sufficient steps toward the task goal to reflect the lack of task-specific knowledge in inexperienced problem solvers: naïve subjects do not know how to define strategies that will take them from the initial state to the goal state. Three examples of preliminary difference-reduction strategies used in the model are Empty (Disk_Stack), Don't-Block strategies and Move (Largest_Disk): Empty (Disk_Stack) seeks to get closer to the task goal by trying to move the currently movable disk of the stack that needs to be transferred to the goal peg. It is a preliminary strategy because it does not specify all the necessary steps to reach the task goal. It also involves an affordance – the property that the stack affords emptying.

Don't-Block strategies reduce differences toward the task goal by avoiding blocking a peg or a disk with another disk. For example, certain strategies in the model avoid blocking the goal peg or the top disk of the stack with the smallest disk 1. Other Don't-Block strategies – called Parity strategies – seek to move the smallest disk 1 optimally when 1 belongs to a structure blocking the top disk of the stack or its destination. Those strategies help empty the stack in relation to the parity of the blocking structure to which 1 belongs. For example, one such strategy blocks the top disk of the stack with Disk 1 when 1 is part of a two-disk structure (Disks 1 and 2) on another peg (e.g., the top disk of the stack is on peg A and 1 and 2 are on B): that move is optimal to avoid blocking the top disk of the stack with Disk 2. The Don't-Block difference-reduction strategies are preliminary: they do not guarantee the successful transfer of the stack to its goal peg. They are based on move affordances – e.g., on the general concern that placing an obstacle at a location corresponding to a destination for an object may restrict the move affordance of that object.

Finally, Move (Largest_Disk) sets the goal to move the largest disk of the stack to the goal peg as a way to get closer to the task goal. It is based on the Largest_Disk affordance mentioned earlier. It emerges from Don't-Block and Empty (Disk_Stack) strategies and refers to the ability to set largest-disk subgoals.

The Critical Role of Empty (Disk_Stack)

Empty (Disk_Stack) represents the main strategy in the early phase of problem solving in the model. It selects the top disk (d_T) of the initial stack to move in priority. As a result, it leads to the acquisition of subgoals that seek to preserve the ability of d_T to move, such as avoiding blocking d_T with another disk in order to empty the stack ('Don't_Block (d_T)" subgoal). Those subgoals constraint the move options offered to the disks occupying the other pegs.

Learning Mechanisms

Observing the outcome of effective or possible moves constitutes an essential learning mechanism in the model. For example, learning that a stack of disks affords emptying occurs by observing that effective moves do empty that stack. Learning an Empty (Disk_Stack) subgoal such as Don't_Block (d_T) takes place by considering possible move options for disks present on other pegs (e.g., in the third state with d_T on peg A, Disk 1 on B, and Disk 2 on C, 1 or 2 can be moved to peg A, but such moves would block d_T.) Other learning mechanisms in the model recode existing knowledge into simpler forms – e.g., simplifying the definition of a series of moves in relation to affordances. The Parity strategy example mentioned earlier – learning to move Disk 1 on top of d_T on peg A when 1 and 2 are on B – illustrates that mechanism. Initially, the model defines the moves of Disks 1 and 2 by assigning a peg for Disk 2 first – based on the relative move affordance of those disks: a peg is chosen such that Disk 2 does not block d_T on peg A. That choice sends 1 over d_T on peg A, but allows the top disk of the stack to move later on, so the sequence of moves for

Disks 2 and 1 is recoded by taking the perspective of the movable disk: In order to empty the initial stack, Disk 1 should block d_T when it belongs to a two-disk blocking structure.

The Acquisition of Largest-Disk Subgoals

The ability to focus on the largest disk and to set largest-disk subgoals emerges from the previous mechanisms. Problem solving begins with a focus on top, movable disks, and evolves – mostly by learning from effective and possible moves. Those mechanisms create the Empty: (Disk_Stack) affordance, resulting in the definition of the Empty (Disk-Stack) difference reduction strategy and of Empty (Disk-Stack) subgoals. A state is then reached where the largest disk of the initial stack stands by itself on the source peg. The model acquires the largest-disk affordance when that disk is moved to the goal peg. It then forms the strategy to move the largest disk to the goal peg as a way to reach the task goal. That strategy, Move (Largest_Disk), now sets priorities for problem solving – taking over Empty (Disk_Stack) and its top-disk focus. After the largest disk of the initial stack is moved to the goal peg, the remaining disks form a new stack on peg B: the largest-disk affordance produces a focus on the bottom-disk of that stack and Move (Largest_Disk) creates a largest-disk subgoal. Table 1 illustrates those mechanisms with a four-disk stack. Problem solving in that example begins with the Move: Movable_Disk affordance and a strategy borrowed from Anzai & Simon (1979)'s participant, that allocates a peg for Disk 1 by avoiding blocking the goal peg with that disk.

Table 1: Acquisition of Largest-Disk Subgoals: An Example With Four Disks.

| State | Disks on Peg | | | Acquired Knowledge |
	A	B	C	
s1	1,2,3,4	-	-	
s2	2,3,4	1	-	[1]AFF (1), [2]DRS (1), [3]S: 1, 2
s3	3,4	1	2	AFF (2), S: 3
s4	3,4	-	1,2	
s5	4	3	1,2	
s6	1,4	3	2	
s7	1, 4	2, 3	-	S: 4, 5
s8	4	1, 2, 3	-	AFF (3), DRS (2)
s9	-	1,2,3	4	

[1]AFF: Affordance (1) Empty:Disk_Stack, (2) Relative_Move, (3): Largest_Disk.

[2]DRS: Difference-Reduction Strategies: (1) Empty (Disk_Stack), (2): Move (Largest_Disk).

[3]S Empty (Disk_Stack) subgoals: (1) Select d_T to move, (2) Move d_T to an empty peg, (3) Don't B lock d_T, (4) Unblock d_T, (5) Block d_T with 1 when 1 and 2 form a two-disk structure on another peg.

The Emergence of Optimal Problem Solving

An interesting property of the above model is its ability to generate optimal solutions to the Tower of Hanoi puzzle. What follows discusses that ability with four and five-disk versions of the game. Two kinds of disk stacks are considered: S_i is the initial stack – the stack of disks present at the beginning state (see Figure 1). S_T refers to a 'transitional' sta ck – a stack that emerges during problem solving and that does not constitute a blocking structure for existing difference-reduction strategies in the model. For example, when the largest disk 5 is moved to peg C, a four-disk stack made of Disks 1, 2, 3, and 4 occupies peg B and is not a blocking structure. In addition, two kinds of largest disks are considered: d_L (S_i), the largest disk of the original stack (e.g., Disk 5 in Figure 1) standing by itself on the source peg, and d_L (S_T), the largest disk of a transitional stack - Disk 4 in the above example.

Optimal problem solving in the model is achieved in two phases which allow the acquisition of – respectively – the ability to focus on largest disks and to set largest-disk subgoals. The first phase ('D isk_Stack Phase') is structured around the initial stack, with a focus on working with movable disks (e.g., the top disk of the stack), using Empty (Disk_Stack) as the main difference-reduction strategy. More precisely, that phase allows optimal moves through affordances – general and task-specific – and preliminary difference-reduction strategies, following the earlier specification of the model: problem solvers begin with a general affordance (Move: Movable-Disk), discover that the stack affords emptying, and use the Empty (Disk_Stack) difference-reduction strategy. The focus on the largest disk of the stack emerges during that phase with inexperienced problem solvers. This is achieved through the mechanisms described in the previous section, which allow optimal moves until the state is reached where the largest disk of the stack is moved to the goal peg – isolating the largest disk as a key disk to move.

The second phase is focused on the largest disks of transitional stacks ('Largest_Disk Phase'). The ability to set largest-disk subgoals – sending the largest disks of a transitional stack to the goal peg – emerges at the beginning of that phase, by means of the processes seen earlier. Problem solving evolves using both largest-disk subgoals and strategies focused on movable disks (e.g., not blocking the top disk of the stack).

Tables 2 and 3 present an overview of the major similarities and difference in optimal problem solving with 4- and 5-disk Tower of Hanoi versions. The mechanisms learning to the acquisition of the Empty (Disk_Stack) and of the Move (Largest_Disk) strategies define the similarities. Problem solving with 4 and 5 disks differs in two ways. Regarding the first move of the smallest disk 1, the optimal strategy with four disks can be characterized as not blocking the goal peg with the smallest disk 1 – Don't -Block (P_G, 1) in Table 2. With 5 disks, Disk 1 moves to the goal peg – 1-to-P_G in Table 3. Moreover, optimal problem solving with 5 disks involves the use of the Parity strategies mentioned earlier – in both the Disk_Stack and the Largest_Disk phases. In that case, optimal problem solving in the second phase is such that the top-disk-driven Empty (Disk_Stack)

strategies are still present and used – if applicable– even if the ability to focus on the largest disk and to set largest-disk subgoals has been acquired.

Table 2: Optimal Problem Solving With 4 Disks.

Issue	Disk_Stack Phase	Largest_Disk Phase
1st Move	Don't Block P_G, 1	N/A
Strategy	Empty (Disk_ Stack)	D_L (S_T) to P_G
Disk Focus	Movable	D_L, Movable
Learning	D_L (S_i) focus	D_L (S_T) focus
	D_L (S_i) to P_G	D_L (S_T) to P_G

Table 3: Optimal Problem Solving With 5 Disks.

Issue	Disk_Stack Phase	Largest_Disk Phase
1st Move	1-to-P_G	N/A
Strategy	Empty (Disk_Stack)	D_L (S_T) to P_G
Disk Focus	Movable	D_L, Movable
Learning	D_L (S_i) focus	D_L (S_T) focus
	D_L (S_i) to P_G	D_L (S_T) to P_G
		Parity Strategies

A Test of the Model's Fit

Anzai & Simon (1979) report a case of almost optimal problem solving with a five-disk version of the Tower of Hanoi: their subject's second attempt at the task (see the verbalizations corresponding to the first ten states in the Appendix.) The authors characterize that episode as implicating 'the goal -peg strategy" (Anzai & Simon, 1979) – a strategy that assigns the successive largest disks of the initial stack to the goal peg. I will now provide an overview of the argument explaining how the present model accounts for the problem solving behavior observed in that episode.

The problem solver in the above study is not entirely novice, her first attempt at the task having resulted in the discovery that the largest disk of the initial stack (Disk 5) should be moved to the goal peg first (C), and that Disk 4 should not go to C. All the moves and verbalizations of that first attempt can be characterized by using the mechanisms from the present model – reflecting affordance-driven behavior motivated by the priority to empty the disk stack (Guimberteau, 2003).

In the second episode, the subject changes her first move, placing Disk 1 on the goal peg because the peg allocation chosen for Disk 1 in the first episode – Disk 1 not blocking the goal peg – did not lead to the goal state: That move is optimal. The affordance-driven mechanisms presented here, augmented with her learning (Disk 5 to Peg C, Disk 4 not to C, Disk 1 to C on first move, and the affordance-related knowledge from her first episode) are able to predict the subject's problem solving behavior. After the first move, her problem solving unfolds through Empty (Disk_Stack): a move of the smallest disk 1 to avoid blocking the open disk of the stack leads the subject to notice that she should have blocked that open disk with 1 instead, to clear a peg occupied by the two smallest disks (fifth state of Episode 2, see lines 31 and 32). She learns here a Parity strategy from

the Don't Block family: with two blocking disks on a peg destined for the open disk of the stack, blocking the open disk with Disk 1 is optimal. Her problem solving behavior develops afterwards through a refinement of that insight. It is first applied in the tenth state to three blocking disks - not blocking the open disk in that case. It is later applied to the goal peg - instead of the open disk of the stack: she blocks the goal peg with Disk 1 that is part of a two-disk blocking structure (state 14), but does not block that peg with Disk 1 belonging to a three-disk blocking structure (state 18).

The present account produces an excellent fit. It predicts the 31 overt moves observed in the episode and 59 of the 62 protocol segments. Two statements not predicted by the model are a comprehension monitoring statement - "I wonder if I've found something new" (line 71 of the episode) – and an activity statement – "This is my way of doing it" (line 74). The third, not predicted, segment is a move questioning – "What?" (line 44) – which is, however, consistent with the affordance-driven behavior proposed by the present model. That statement occurs after a move of the smallest disk to the goal peg: there should not be any questioning of that move if it is driven by one of the Parity strategies from the model. It is likely that the subject – trying to remove the two smallest disks 1 and 2 to unblock the largest disk to move it to the goal peg – is focusing on Disk 2 first because it is the larger of the two blocking disks. She has explicitly exhibited that behavior in her first problem solving episode. That preoccupation with Disk 2 is consistent with the present model: it is based on the relative move affordance seen earlier – the fact that a larger disk should be considered first because it affords moving less than a smaller disk. Moreover, that focus on Disk 2, not on Disk 1, can be reconciled with a top-disk account by saying that the subject infers a blocking pattern for the smallest disk – the same way she has been using Parity strategies so far: moving Disk 1 such that it blocks the goal peg is optimal when Disk 1 is part of a two-disk blocking structure on another peg.

We can estimate the model's fit by the total number of moves and verbal segments accounted for divided by the total number of observed moves and verbal segments, following similar practices (e.g., Newell & Simon, 1972; VanLehn, 1991). Using that calculation, the analysis fits [(31 − 0) + (62 − 3)] / 93 or 96% of the protocol – an excellent outcome.

## A New Class of Optimal Tower of Hanoi Strategies

A comparison of the present model to existing optimal Tower of Hanoi strategies from the cognitive science literature (Simon, 1975) reveals that the model offers a new class of problem solving strategies. Three classes of optimal strategies have been identified so far (Simon, 1975): Goal-Recursion, Perceptual (Basic and Sophisticated) and Move-Pattern. Goal-Recursion is the recursive subgoaling strategy mentioned earlier that sets largest-disk subgoals (Altmann & Trafton, 2002; Anderson & Lebiere, 1998; Simon, 1975). The Basic Perceptual strategy (Simon, 1975) identifies the largest disk d_L not yet on the goal peg,

emptying up first the source peg, then the goal peg, from disks smaller than d_L. The Sophisticated Perceptual strategy (Simon, 1975) clears the largest blocking disk from both the source and the goal pegs until no such blocking disk exists to move d_L. The Move-Pattern strategy is based on cycling patterns – e.g., relying on a pattern between the moves of the smallest disk and the parity of the moves, and on a peg assignment cycle for the smallest disk in relation to the parity of the total number of pegs! None of the above classes are structured around affordances. Two of those classes – Goal-Recursion and Perceptual – are based on means-ends analysis, but those difference-reduction strategies are not affordance-driven. In other words, the present model embodies a new class of problem solving strategies for the Tower of Hanoi task.

Skilled Learning

One virtue of this new model lies in its capacity to shed an entirely new light on the case of skilled learning displayed by Anzai & Simon (1979)'s participant. The subject's four strategies – Selective Search, Goal-Peg, Recursive Disk-Subgoaling and Pyramid-Subgoaling – can be re-characterized through affordance-driven learning, with excellent fits (Guimberteau, 2003). That capacity is well illustrated by the example of the subject's third strategy: Recursive Disk-Subgoaling plans the move of each disk of the initial stack, beginning with the largest at the bottom and continuing with the next-largest until the top disk is reached (see Figure 2). That strategy and the mechanisms behind its discovery can be described in the present model without invoking recursion: the subject moves up the initial stack of disks to select disks to move, and she assigns pegs by repeating the application of a move relation that has been recoded from a blocking relation. The discovery involves looking for ways to use the stack of disks by moving up, instead of moving down – a defining characteristic of the Empty (Disk_Stack) strategy applied so far.

The above example emphasizes that successive subgoaling – subgoals applying to successively smaller disks – does not necessarily reflect recursive subgoaling – subgoals resulting from a strategy that decomposes a problem into smaller versions of itself. In that respect, the model offers valuable insights for common research practices in cognitive science that tend to infer recursive strategies based on the observation of successive subgoals.

79 So, if there were three… yes now it gets difficult.
80 Yes, it's not that easy….
81 … This time, 1 will…
82 Oh, yeah, 3 will have to go to C first.
83 For that, 2 will have to go to B.
84 For that, um… 1 will go to C.

Figure 2: The Recursive Subgoaling Strategy: First Instance (Anzai & Simon, 1979).

Discussion

The model described here proposes that a central aspect of Tower of Hanoi problem solving – the ability to focus on and to use the largest disk of a pyramid – assumed to be elementary in most existing models (e.g., Karat, 1982) – is structured around the acquisition and use of task-specific affordances. An implication of this characterization is that an explanation of strategy change in inexperienced problem solvers needs to consider the above mechanisms.

The fact that the model is able to develop problem solving competence through affordance-driven mechanisms, starting from simple Empty (Disk_Stack) and Don't_Block strategies, is noteworthy as well. Those new mechanisms provide fresh insights regarding learning and strategy acquisition in Anzai & Simon (1979)'s protocol (Guimberteau, 2003). This suggests that they have the potential to shed new light on the nature and the origin of strategies discovered by problem solvers. Two remaining tasks regarding the development of this model are its implementation and its experimental study. I am currently working on those two issues.

The present analysis offers a process model of Tower of Hanoi problem solving based on affordances. Previous research has emphasized the essential role played by external representations in cognition (e.g., Zhang & Norman, 1994) – leaving out the question of affordance-related mechanisms that underlie problem solving with that task. The processes put forward in the research can be used to examine the distributed nature of learning and problem solving in other task as well.

In summary, the model proposed here suggests that the ability to focus on the largest disk of a pyramid during problem solving with the Tower of Hanoi task develops from affordance-driven strategies. Moreover, it is possible to identify a new class of optimal strategies based on the above model – accounting for the development of problem solving competence at early stages with the Tower of Hanoi task, and offering insights regarding the development of more complex strategies. The new class of affordance-driven strategies identified in this research is valuable to help shed new light on learning and problem solving.

Acknowledgments

References

Altmann, E.M., & Trafton, J.G. (2002). Memory for goals: An activation-based model. *Cognitive Science*, *26*, 39-83.

Anderson, J.R. & Lebiere, C. (1998). *The atomic components of thought*. Hillsdale, NJ: Erlbaum

Anderson, J.R., Kushmerick, N. & Lebiere, C. (1993). The Tower of Hanoi and goal structures. In J.R. Anderson (Ed.), *Rules of the mind*. Hillsdale, NJ: Erlbaum.

Anzai, Y., & Simon, H. A. (1979). The theory of learning by doing. *Psychological Review*, *86*, 124-140.

Guimberteau, B. (2003). *Designing theories of learning: New insights*. Ph.D. Dissertation, University of California, Berkeley, CA.

Karat, J. (1982). A model of problem solving with incomplete constraint knowledge. *Cognitive Psychology*, *14*, 538-559.

Newell, A., & Simon, H. A. (1972). *Human problem solving*. Englewood Cliffs, NJ: Prentice-Hall.

Norman, D. A. (1988). *The psychology of everyday things*. New York: Basic Books.

Norman, D. A. (1993). *Things that make us smart*. Reading, MA: Addison-Wesley.

Ruiz, D. & Newell, A. (1989). Tower-noticing triggers strategy-change in the Tower of Hanoi: a Soar model. *Proceedings of the Eleventh Annual Conference of the Cognitive Science Society (pp. 522-529)*. Hillsdale, NJ: Erlbaum

Simon, H.A. (1975). The functional equivalence of problem solving skills. *Cognitive Psychology*, 7, 268-288.

VanLehn, K. (1991). Rule acquisition events in the discovery of problem solving strategies. *Cognitive Science*, *15*, 1-47.

Zhang, J., & Norman, D. A. (1994). Representations in distributed cognitive tasks. *Cognitive Science*, *18*, 87-122.

Appendix – Second Problem Solving Episode of Anzai & Simon (1979) – Verbalizations Corresponding to the First Ten States

25. Let's see…I don't think 5 will move.
26. Therefore, since 1 is the only disk I can move, and last I moved it to B, I'll put it on C this time … from A to C.
27. So naturally, 2 will have to go from A to B.
28. And this time too, I'll place 1 from C to B.
29. I'll place 3 from A to C.
30. And so I'll place 1 from B..to C.
31. Oh, yeah! I have to place it on C.
32. Disk 2..no, not 2, but I placed 1 from B to C..Right?
33. Oh, I'll place 1 from B to A.
34. (Go ahead.)
35. Because..I want 4 on B, and if I had placed 1 on C from B, it wouldn't have been able to move.
36. 2 will go from B to C.
37. 1 will go from A to C.
38. And so, B will be open, and 4 will go from A to B.
39. So then, this time…It's coming out pretty well…