The Pennsylvania State University

The Graduate School

THE ORIGINS AND DEVELOPMENT OF DOMAIN-GENERAL PROBLEM SOLVING STRATEGIES

A Thesis in

Psychology

by

Katherine H. Grobman

Submitted in Partial Fulfillment

of the Requirements

for the degree of

Master of Science

August 2002

We approve the thesis of Katherine H. Grobman.

Date of Signature

Rick O. Gilmore
Assistant Professor of Psychology
Thesis Advisor

Lynn S. Liben
Distinguished Professor of Psychology

Paul A. Klaczynski
Associate Professor of Psychology

Keith Crnic
Professor of Psychology
Head of Department of Psychology

Katherine H. Grobman

Abstract

Despite the central role of problem solving in cognition and despite the extensive study of problem solving in young children, we do not have a clear answer to the basic developmental question of the origins of broadly applicable "domain general" problem solving strategies. Do these strategies emerge gradually or all at once? Are these strategies innate or are they learned? This experiment teases apart some domain general strategies by presenting infants two types of toy-retrieval problems that differ only in which strategies may solve the problem. In each problem, a toy was placed on a cloth that extended through the entire length of a canal. When infants pulled on the cloth in the obtuse-canal problem, the toy continually moved closer to them. When they pulled the cloth in the acute-canal problem, the toy moved further away before moving closer to them. The obtuse-canal problem is solvable by hill-climbing or means-ends analysis. The acute-canal problem is solvable only by means-ends analysis, because it involves moving further from the goal. Results show that 9 months olds could not use either strategy and 16 month olds were comparatively better at using both strategies. The 12 month old infants could use hill-climbing just as well as the 16 month old infants but they were no better at using means-ends analysis than the 9 month old infants. This suggests that domain general problem solving strategies develop gradually throughout infancy. Twelve month old infants showed that with experience, they were able to solve the problem as well as 16 month olds. This suggests means-ends analysis can be learned within a particular context.

Table of Contents

List of Figures

List of Tables

Acknowledgements

So many wonderful people helped make this research possible; it is a privilege to get to acknowledge some of them here. Rick Gilmore, I will never forget our spirited discussions and your persistence even when I became overwhelmed. Thank you. Lisa Szechter, thank you for your unwavering love and support no matter how much rain falls in our lives. The way you brighten my life gave me the hope I needed to complete this work. Lynn Liben, thank you for believing in me and for your thoughtful suggestions that brought to my attention literatures I may not have considered. Ulrich Mueller and Paul Klaczynski, thank you for your feedback on drafts of my thesis. Mom and Dad (Sondra & Gil Grobman), thank you for all of your help that is too intangible to put into words. And, dad, thank you for putting your very tangible carpentry expertise into the experimental apparatus. Many members of Rick Gilmore's lab helped make this project possible by recruiting participants and coding data: Melissa Arnett, Tom Baker, Jen Bealer, Liane Cole, Thea Gutteridge, Alisson Jazwinski, Brittany Nogales, and Katherine Watson. Thank you. I would also like to thank Sigma Xi, the Scientific Research Society, for a small research grant in aid of this project. Finally, I would like to thank all of the infants and parents who volunteered to participate in this research. Your time and altruism was the foundation for everything found within the pages of this thesis.

The Origins of Problem Solving and the Development of Domain-General Strategies

A number of psychological researchers have understood problem solving as a framework for understanding higher-level cognition (Newell, 1980), and as a bridge to unify our understanding of learning and performance (Anderson, 1993). Although the development of problem solving has been studied extensively in young children (e.g. Alibali, 1999; Gardener & Rogoff, 1990; Haith, 1994; Karmiloff-Smith, 1986; Schauble, 1990; Siegler, 1976), we do not yet have clear answers to the basic developmental question of how broadly applicable problem solving strategies originate in infancy. The present study is concerned with how infants learn strategies for particular tasks, and how they use their domain general problem solving abilities when a situation is unfamiliar. This study was designed to address the developmental question of the origins of children's problem solving strategies, and to examine the possible mechanisms underlying the development of these strategies.

Information Processing Framework

Among the most powerful tools for studying problem solving is the information processing framework. The information-processing framework has been used to study problem solving from childhood to adulthood and across many fields such as Psychology and Computer Science. One strength of this framework is its precise way of defining problem solving situations.

A *problem* is any situation in which a person, or some other agent, has a goal that has not yet been achieved. An obstacle or obstacles needs to overcome to achieve that goal. Goals can be thought of in very general terms to include anything an agent might want. Goals can include a diverse range of objectives such as finding values for “x” which make an algebraic sentence true, and satisfying hunger. Obstacles can also be thought of in very general terms. Some obstacles can be dramatic and fulfill our conventional definitions of obstacle. Other obstacles can be so minor we hardly notice them. For example, if an adult is asked to multiply 6 by 8, he can probably say “48” without much effort. In this case, the obstacle to overcome was simply the retrieval of a fact from memory. Yet if a child is asked the same problem, she might have a more significant obstacle to overcome. In this sense, we could describe life as problem solving and even characterize some of development as growing from a novice to an expert at living life.

Newell (1982) suggested that all rational activity can be described, on the level of mental phenomena, in terms of problem solving. He introduced terms that have evolved over the decades into a more mathematical and concrete terminology. Newell described “states of knowledge” and this has become a “problem space” as implemented in computational models such as Soar (Michon & Akyuerek, 1992) and STRIPS (Filks & Nilsson, 1971). A *problem space* is a set of all possible state-of-affairs/knowledge and their inter-relationships.

Figure 1.1. A problem space for a Luchins Water-Jug Problem

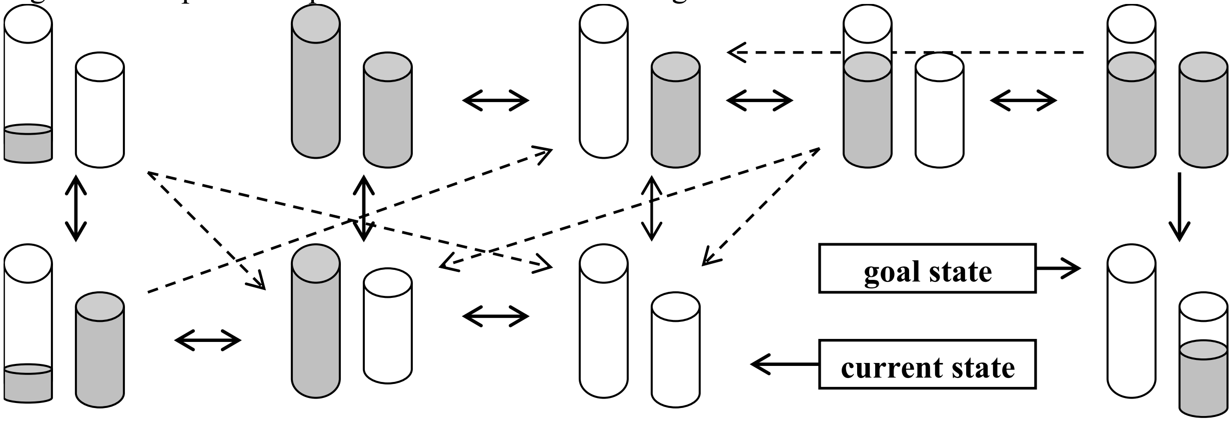

To illustrate a problem space, and other formal terminology, consider a sample problem based on the classic Luchins (1942) Water-Jug problems. Suppose we have only an empty three-ounce bottle, an empty four-ounce bottle, and a large pool of water. We want exactly 2 ounces of water in the 3-ounce bottle but neither bottle is marked. The *current state* represents the state of affairs at this moment in problem. When we begin, the current state is the *start state*, which is two empty bottles. The *goal state* is a way of representing the person's goal. In this case, it is two ounces of water. *Operators* are ways we can change our current state into a new state. In this example, the operators are filling a bottle with water, emptying a bottle of war, and pouring water from one bottle to the other. Figure 1.1 illustrates the problem space with the current state and goal state denoted. Arrows represent the operators. A problem is considered *solved* when the current state is the goal state.

Problem Space: A Personal, Theoretical, & Social Construction

The use of a problem space provides a useful theoretical construct for understanding domain general problem solving. It provides a language for describing the essential aspects of any problem in a formal way and consequently allows us to compare and contrast seemingly different problems by varying some elements and not others. With this formal language, we can compare and contrast problems as seemingly different as perplexed physics students' first attempts to do their homework and infants' first efforts to get toys.

The formal language of this approach to problem solving may make it appear that problems are objective representations. In some ways, this is certainly true. Yet, there are also elements of the problem space, which are intimately tied to social, personal, or theoretical constructions. For example, how a problem's social context (cover story) can make a problem hard or easy has been examined with problems whose problem spaces are structurally equivalent (isomorphic) but which have different cover stories (e.g. Kotovsky & Simon, 1990). This example shows how the social construction of the problem space is not necessarily a weakness of the model. It is a component of problem solving. The goal-state is also something open to interpretation. Just because an experimenter assigns a particular goal, this does not mean the problem-solver has adopted that goal. The components of the problem space state are also open to interpretation. For example, in the problem space I constructed of a Lunchins Water-Jug Problem (Figure 1.1), I did not represent the large pool of water even though the amount of water in it changes. I did not consider it a significant enough part of the states of affairs to represent.

Similarly, I did not include the infinite intermediate steps of pouring water where a jug is 100% full, 90% full, 80% full, and so on. The problem space is also a personal construction of the problem-solver. Indeed, some researchers in Artificial Intelligence have gone further by distinguishing between a "problem space" and "state space," where the latter is the complete search space and the former is the representation in the actual problem-solver's mind. This distinction is not made in some literatures, such as scientific discovery & development literature (e.g. Klahr & Dunbar, 1988). I will not be using this distinction because, unlike those designing artificial intelligence, psychologists can never distinguish what a person demonstrates knowledge of (the state space) and what a person's actual knowledge might include (i.e. more parts of the problem space than demonstrated).

Domain General versus Domain Specific

Given this formal language, we can describe problem solving by the way an agent traverses the problem space. A problem-solver may have skills for traversing particular types of problem spaces. For example, consider an algebra student encountering a particularly difficult algebra problem. She can approach solving the problem with techniques learned especially for algebra such as "do the same thing to both sides." Techniques used especially to traverse particular types of problem spaces are referred to as *domain specific* problem solving strategies. There are occasions where a problem solver has no techniques specifically for the domain of a problem. Instead, the problem solving is *domain general*, involving only those properties which all problems share. Any problem can be thought of in domain-general and domain-specific ways. The same algebra problem could be approached without background knowledge of the domain. The student may use the domain general strategy of "trial and error" by "guessing numbers for the variable and testing if it works." It is in this way that the same formal language can be used to describe a person's development from novice to expert. Given this development from novice to expert, the origins of problem solving can be studied by asking where domain general problem solving strategies come from.

Though we enter a new area with domain general strategies, we can develop more efficient domain specific strategies through experiences such as instruction and through the development of our conceptual understanding of the domain (e.g., Rittle-Johnson & Alibali, 1999; Grobman, 1998). While there are many general approaches to problems such as planning, brainstorming, and clarifying the problem, I would like to consider domain-general problem solving *strategies*, domain-general techniques for traversing a problem-space. There are four established domain-general problem solving strategies: analogy, trial-and-error, hill-climbing, and means-ends analysis.

Trial-and-Error: a Focus on Action

Trial-and-error is a problem solving strategy that involves applying randomly chosen operators to get from the current state to the goal state. This strategy requires very little background knowledge because it focuses on generating different actions to try to solve the problem. Despite the simplicity of this strategy, it is sometimes a challenging strategy to use because the operators are difficult for the problem solver to identify or apply. Examples include many types of creative problem solving situations such as divergent thinking tasks like the nine-dot problem (Burham & Davis, 1970), or functional fixedness problems like Dunker's candle problem (Dunker, 1945). Trial-and-error may also be a difficult strategy to use when there are too many possibilities for solving the problem. For example, in the proverbial problem of "trying to find a needle in a haystack" a trial-and-error strategy is simple: keep applying the operator "look at a little bit of hay." When the search space is large, the task can easily become overwhelming however, in these circumstances, it becomes clear that it is grossly inefficient to solve a problem by trial and error.

Hill-climbing: a Focus on Difference-Reduction

Hill-climbing is a problem-solving strategy that involves choosing an operator or operators to reduce the discrepancy between the current state and the goal state. The maxim of a "journey of a thousand leagues begins with a single step" illustrates how effective hill-climbing can be in solving overwhelmingly large problems. Rather than considering a single operator to get a thousand leagues away, the problem solver takes a single step to adjust the current state to bring him slightly closer to the goal state. The operator "taking a single step" can then be applied repeatedly to each new current state to reduce the discrepancy between the current state and the goal state (distance to destination). Hill-climbing is a considerably more efficient strategy to complete a journey than trial-and-error. Using a trial-and-error strategy, one would repeatedly apply the operator "take a single step" randomly. Although a traveler could eventually reach his destination by walking randomly, always walking in a direction that gets him closer to his destination makes the journey shorter. The measure of ‘distance’ to the goal does not need to be a metric of length. For example, solving a Rubik’s cube puzzle requires making each of the six sides a single color. In this case, having three sides as a single color is closer to the goal of six sides than two sides of a single color.

Means-Ends Analysis: a Focus on Causal Relationships

Means-ends analysis is a strategy of solving a problem by breaking it down into smaller problems. That is, sometimes when we have a problem, we recognize that a particular operator would solve the problem from a particular state. However, we might not be able to apply the desired operator until we reach state by solving another problem. Solving this second problem becomes a *sub-goal*. Means-ends analysis involves solving a problem by one or more recursive sub-goals. If your goal is to write a letter, you first need to get a pen and paper. Getting the pen and paper is a sub-goal. In order to use this strategy, you need to recognize causal relationships. The larger problem has a means to an end. The smaller problem has a means to an end. And the smaller problem’s ends is the larger problem’s means.

Analogy: a Focus on Similarities among Different Contexts

Analogy is unique as a domain-general strategy in that it allows a problem-solver to transfer a domain-specific strategy from one domain to another. Analogy involves a *base* problem with a known solution and a *target* problem, which has not been solved (Goswami, 1996). To use analogy you must recognize a similarity between the target and base problems. A strategy comparable to that used in the base problem may then be used to solve the target problem. For example, you discover while cooking that a food is too salty but you cannot take the salt out. To make the food less salty, you put more of the other ingredients into the food. Later that day you start an acrylic painting class. You discover while painting that you mixed too much blue into a paint mix but you cannot take out the blue. How do you make the paint less blue? You might recognize how similar the paint and food problems are: elements are mixed, once mixed an element can not be removed, and the problem is having too much of one element. Even though you know little about painting, you might try the same operators you used when solving the base problem of cooking to solve the target problem of painting. A solution could be to add more of the other colors to your paint mix.

The Development of Problem Solving in Infancy

Though their experiences may be quite different from ours, infants, like adults, confront problems in their everyday lives. Moreover, like adults, infants are able to solve some of their problems. Henderson and Dias (1987) observed four infants from 4 to 18 months of age in their homes for 75 waking hours. The infants were observed encountering a problem once every three minutes. Most of the infants' problems were social (50%), such as wanting to get mom's attention or physical (39%), such as retrieving a distant toy. Infants solved about two-thirds of problems, and most problems were solved by the infant's direct action or compliance with a social demand. Older infants were more likely than younger infants to solve problems on their own. In addition to forming a preponderance of infants' problem solving experiences, social problem solving reveals some of infants' earliest problem solving competencies. From a socio-cultural perspective, Mosier and Rogoff (1994) showed how infants increasingly engage in "guided participation" to instrumentally use their mothers in order to solve problems.

Clearly infants solve problems but how does infant problem solving *develop*? In particular, where do domain general problem solving strategies come from? The most extensively studied aspect of infant problem-solving may be analogical transfer. Yet analogy, by its nature, can not address origins of domain-general problem solving strategies. Piaget provides the most extensive account of development of problem solving. Indeed, the research since Piaget has largely been a refinement, or extension to his theory.

Infants' Use of Analogy to Solve Ever-More Domain-Specific Problems

Perhaps infants develop their problem solving skills by analogical transfer. At first, this may seem implausible, because even adults have difficulty with analogical transfer. A classic example is Dunker's (1945) "Ray and Tumor" problem, in which participants are presented with a challenging problem, and immediately afterward they are asked a structurally identical problem with different surface features. Adults, like children in comparable studies, rarely notice the commonality (Brown, 1989). Children rarely show transfer in laboratory studies unless similarities across problems are highlighted: either by explicitly telling them about the similarities or by using the same colors, shapes, or labels (e.g., Brown & Campione, 1984). However, if the operators (e.g., pulling, pushing) embedded in the problems are the same, 3- to 4-year-old children can transfer strategies (Brown & Kane, 1988). Brown (1990) observed children from 17 to 36 months attempt to retrieve a toy by choosing the appropriate tool from an array. Younger children needed a parent for the first task where, for example, a hook was used to bring the toy closer. However, there was no age difference in the ability to transfer this strategy to an analogous problem where, for example, a rake brought the toy closer. Over 90% of the children could transfer this strategy. Even infants as young as ten months of age have been shown to use analogical transfer to solve problems when the means rely on the same causal mechanism (Chen, Sanchez, & Campbell, 1997). In a microgenetic study which closely resembles Brown's (1990) tool use experiments, Chen and Siegler (2000) suggested that the same components and the same processes that underlie adult and older children's problem solving development, also explain 1 to 2 year old infants' problem solving strategies.

Analogy is a fascinating domain general strategy to study in the development of children’s problem solving, because of the far-reaching implications for children’s use of analogy. If children use analogy in their everyday experience, we can construct models of cognitive development that do not depend on children’s knowledge of universally valid principles (Goswami, 1996). For example, do children acquire Piaget's concept of conservation as an abstract principle and then apply it to all contexts or do they acquire conservation in a single context and transfer that knowledge to new contexts by using analogies? Yet precisely because problem solving by analogy requires previous domain specific problem solving strategies, its use does not give us particular insight into the origins of domain general problem solving strategies.

Piaget’s Theory: Differentiation of the Means & Ends

One possible way to understand infant’s problem solving development is to examine an infant's understanding of causality. Piaget (1954) provides a detailed description of the infant's developing understanding of causality through the gradual differentiation of different aspects of a problem solving situation. Consider one of Piaget's observations that demonstrates a more complete causal understanding:

> Since 1:0 (10) Jacqueline has known how to utilize the laws of swinging. Seated

> in her baby swing she gives it an increasing momentum, then turns backward, raises her legs and lets herself swing, remaining perfectly quiet until the movement stops. The difference between this attitude and the attitudes that have previously characterized her play is this: up to now she has been constantly active in her swing, as though her movements were necessary to make the phenomenon last, whereas now she knows that the action itself is controlled by and can be referred to the laws of the swing (Piaget, 1954, p.288).

From birth through Piaget's sub-stage 3, infants do not distinguish their own actions ("efficacy") from the means ("phenomenalism") or the ends. All of these states are viewed as one indivisible whole. Stage 4 infants distinguish the ends but still understand the means as equivalent to their own actions. At Stage 5, infants distinguish all three aspects of the problem solving situation. In the above example, Jacqueline is at stage 5 because she recognizes that she can swing her legs (her own actions), that she can start the swing moving by swinging her legs (a means), and that she wants the swing to move (ends).

Extensions of Piagetian Theory: Intentionality & Goal-Directedness

Experimental evidence suggests that infants may become increasingly able to express intentions in their problem solving. Operational definitions of “intentionality” examine the infant's attention to the goal and to the infant's readiness to use operators, often called the "means." One example of increasing intentionality comes from Adele Diamond's Object Retrieval Task (Diamond, 1994). In this task a toy is placed inside a transparent box with only one open side. The box is given to an infant with the open side not facing the infant. From 6 to 12 months of age, infants become more capable of inhibiting direct action for their goal (toy) from their means to the goal (reaching & turning the box).

In one longitudinal study, 6 to 8 month old infants attempted to retrieve a toy that was placed behind a cup (Willatts, 1984b). Infants' intentionality was scored from 0 (no intentionality) to 3 (high intentionality) based on their initial behavior toward the cup. Specifically, infants were judged to be intentionally solving the problem if they lifted the cup and released it, or lifted the cup and held it passively. Infants increased their intentional problem solving with age according to this scale. As infants get older, they become capable of retrieving a toy on a longer and longer cloth (Willatts, 1999). Longer cloths are more difficult because an *initial* pull (means) does not move the toy far enough to retrieve the toy (ends). Infants must maintain their focus on achieving their goal to retrieve the toy. Doing so demonstrates a gradual increase in infants' intentional use of procedures towards the goal.

Bates, Carlson-Luden, and Bretherton (1980) also coded infants' intentionality during tool use problem solving, by coding behaviors directed towards the means, and behaviors directed towards the ends. That is, how much of the infant's attention was focused on the operator (in this case, tools are the means) and how much of infant's attention was focused on the goal (in this case, the toy is the ends). Among nine to 10-month-olds, Bates et al. (1980) found that intentionality varied from problem to problem. Because such a narrow age ranged was studied, age-related development could not be assessed. However, their finding suggests that intentionality may not be a quality that changes on a global level, but rather within specific contexts.

Refinement of Piagetian Theory: the addition of Information Processing

Case (1985) proposed a Neo-Piagetian theory that implicitly assumes the use of sub-goals by infants. He conceptualizes problem solving in his model as the core mechanism underlying cognitive development through stages. Case experimentally demonstrated that 8 to 12 month old infants could separate the "means" from the "end" by solving a bell-balance beam problem (Figure 1.2).

Figure 1.2: Adapted from Case (1984). An example of a problem solving task, and the process Case implicitly assumes infants use to solve the problem.

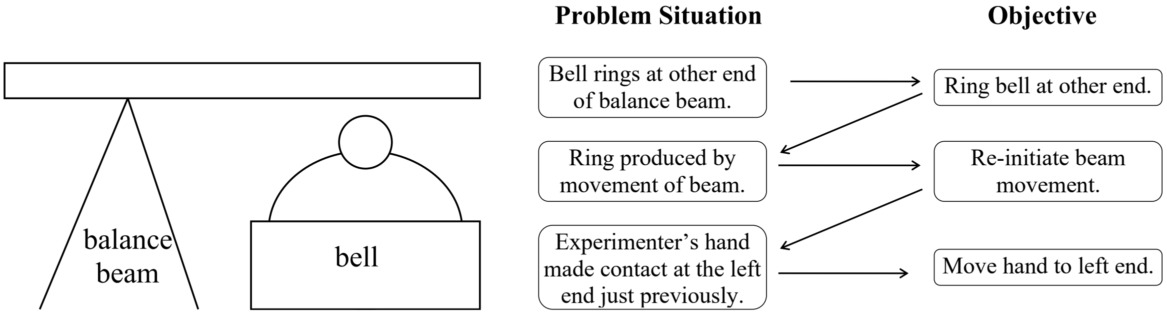

Throughout his book showing problem solving and Piagetian stages across development, Case (1984) illustrates how problem situations lead to objectives which lead to further problem situations and further objectives. According to figure 1.2, this is no different for infants. After an experimenter rings the bell by pushing the balance beam up, the infant attempts to produce the same sound. According to the diagram, infants first adopt the goal (ring bell), recognize the means used by an experimenter (move beam), and then set a sub-goal (move beam). Based on their observation of the experimenter, the infants set a new sub-goal (to achieve the first sub-goal) of moving their hand. The actual process of setting sub-goals (means-ends analysis) is not observed during the experiment. Rather, it is a process assumed to underlie the infant's response.

The implicitly assumed use of means-ends analysis becomes explicit when we consider further developmental theories that construe the child as a 'universal novice' and development as the acquisition of expertise (Wellman & Gelman, 1998). Some cognitive psychologists study the formation of expertise and skill acquisition as aspects of problem solving (e.g., Chase & Simon, 1973). Anderson's ACT* computational model provides a prominent example to illustrate these learning processes (Anderson, 1993). This model assumes that problem-solving is governed by a means-ends problem-solving structure, where learning is a complex interaction of explicit declarative knowledge and procedural problem-solving abilities (Anderson, 1995). When encountering a difficult problem, the model searches for a similar previously solved problem and attempts to solve the new problem by analogy (Anderson, 1993; Anderson, 1997). Anderson (1993) has stated that based on primate research, his "own belief is that the means-ends problem-solving method is an innate part of cognitive machinery of humans and other primates." This assumption is understandable because means-ends analysis acts like an 'engine' within the ACT* model.

According to Anderson (1989), a “litmus test” for the adequacy of a theory of knowledge is that something meaningful be said about how knowledge falls into three philosophical categories of knowledge: nativism (innate), empiricism (experience), and rationalism (constructed). Anderson shows how this litmus test is met by the PUPS theory, which is a step in his modeling development between ACT* and ACT-R. Weak-methods, including means-ends analysis, are a built-in part of knowledge according to Anderson's models. Anderson and others are able to model many aspects of human cognition including problem solving. However, to my knowledge, none of the aspects of the ACT models require that means-ends analysis abilities always be present. The development of expertise and skill acquisition does not necessarily depend on the same developmental process that underlies the formation of domain general strategies. Using the 'engine' analogy, ACT models tell us what the engine does but not how or when the engine was put together.

Do Domain General Problem Solving Strategies Develop?

Previous research on the development of domain-general problem solving strategies has focused on analogical transfer. As previously discussed, studying analogy presupposes the ability to solve some problems already. This is consistent with the implicit and explicit assumption of the information-processing approaches. Infants possess a complete repertoire of domain general problem solving strategies. That is, these strategies simply *are*; they do not develop. Yet, paradoxically, the research concerned specifically with infant's problem solving suggests that there really is development of problem solving throughout infancy. In particular, the Piagetian approach of gradually differentiating the means and ends most closely resembles domain general problem solving strategies. That is, information-processing "operators" are like the Piagetian "means" and the information processing "goal-states" are like Piagetian "ends."

Experimental Evidence Contradicting Piagetian Theory

One possible way to resolve this paradox is to question Piaget's observation of development. Decades of research by Willatt's (e.g. 1984a, 1999) illustrate this approach with a prototypical experiment of physical problem solving by infants. His experiments involve enticing an infant with a toy, placing the toy out of reach, and providing the infant with some means of getting the toy. A number of factors influence the age at which infants can successfully retrieve the toy. In one experiment (Willatts 1984a), two cloths were laid in front of the infant with the toy placed on one cloth. Nine month old infants could retrieve the toy, but 6 month old infants pulled at random. In another experiment (Willatts 1984a), a barrier was placed over the edge of a single cloth. When the toy was placed at the other end of the cloth, 9 month old infants removed the barrier and pulled the cloth to retrieve the toy. However, in a control condition where the toy was placed beside the cloth, 9 month old infants played with the barrier. Willatts (1999) has suggested that infants solve these types of problems through sub-goals, by means-ends analysis. Despite this empirical evidence, further evidence may complicate Willatt's conclusion. In an experiment with a similar design, a small curtain hung between the infant and the toy (Munakata, McClelland, Johnson, & Siegler, 1997). The curtain was either transparent or opaque. Seven month old infants could retrieve the toy if they had to pull it under a transparent curtain. They could *not* retrieve the toy from behind an opaque curtain. Munakata et al. (1997) suggest that this finding indicates that errors on the A-not-B task are due to conceptual difficulties with object permanence and not simply search errors. More importantly for this discussion, another possibility they acknowledge is that infants fail to pull the cloth through the opaque barrier because they do not have initial feedback that the toy is moving closer. That is, perhaps this prototypical experiment of infant's problem solving fails to clearly distinguish between the use of means-ends analysis and hill-climbing.

Theoretical & Operational Differences between Hill-Climbing & Means-Ends Analysis

Means-ends analysis and hill-climbing may seem conceptually similar. In some work, means-ends analysis is sometimes used as a broad term to include both setting sub-goals and difference-reduction (e.g. Polson & Jeffries, 1982). This simplification is understandable from a formal standpoint. Afterall, we might imagine someone saying means-ends analysis is a special case of hill-climbing. This would make means-ends analysis a hill-climbing strategy where the organization of sub-goals becomes the metric for reducing discrepancy. Alternatively, we might imagine someone saying that hill-climbing is a special case of means-ends analysis. Difference reduction would be understood as successive resolution of sub-goals. But does a formal equivalence mean there is an empirical equivalence? To elaborate on this question, consider the following prototypical problem-solving situation as an illustration: An infant sits in her father's lap. A toy star, which she wants, is placed out of her reach upon a cloth that stretches to her. How does she get the toy? Figure 1.3 describes the problem space.

Figure 1.3 Sample problem-space of a prototypical object-on-cloth infant problem solving task

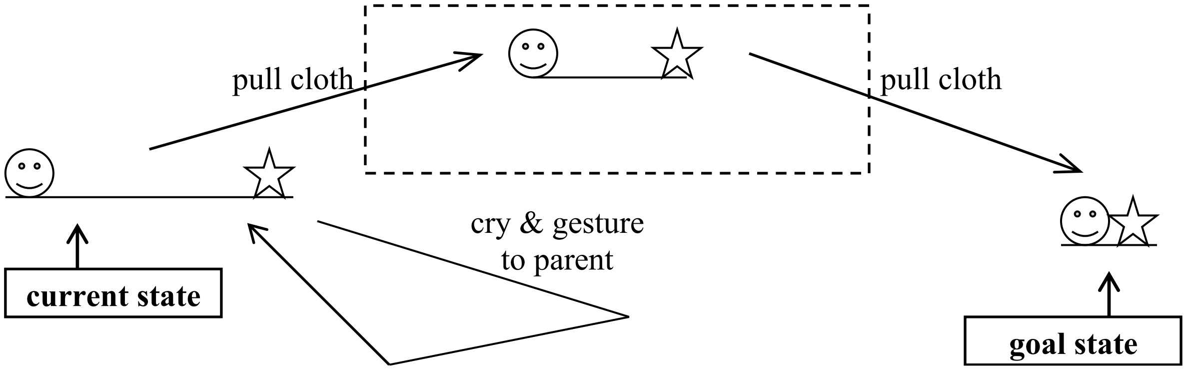

The infant might successfully retrieve the toy star by using several different domain general strategies; consider three strategies in turn. First, consider trial-and-error. She might try gesturing toward the toy and her father. This does not work. She then tries pulling on the cloth. This moves the toy but she still does not have the toy. She then might try crying. This does not work. She tries pulling the cloth again. This time the toy reaches her, and she has retrieved the toy. To use trial and error, she focused on her own actions and could determine her success by seeing if she got the toy. Second, consider hill-climbing and how she pulls the cloth. This provides her with feedback (illustrated by the intermediate state in the dashed box). She notices that the toy is closer. That is, the distance between her current state and goal state is reduced. Given this affirming feedback, she pulls the cloth again and retrieves the toy. To use hill-climbing, she focused on the feedback she got about difference reduction. Third, consider that she might use means-ends analysis and pull the cloth. This provides her with feedback (illustrated by the intermediate state in the dashed box). She notices that pulling the cloth is a means to an end of moving the toy. She notices that getting the cloth would mean getting the toy. That is, she can consider getting the cloth as a sub-goal on the way to getting the toy. Given this affirming feedback, she pulls the cloth again and retrieves the toy. To use means-ends analysis, she focused on the feedback she got about causal relations.

Notice how each strategy was successful by traversing the same problem-space. In particular, both means-ends analysis and hill-climbing hinged on the feedback the infant took from the intermediate state. This illustrates the appeal of reducing these strategies to a single strategy. Nevertheless, the sort of feedback the infant took from the intermediate stage is different for each strategy. This suggests that, although hill-climbing and means-ends analysis could be considered formally equivalent, they are not equivalent in the experiences of problem-solvers. Can this difference in the internal interpretations of problem-solvers be studied empirically? Since the internal differences between the strategies hinge on feedback, we could empirically separate the strategies by designing a task that would give positive feedback when using one strategy and negative feedback when using the other strategy. The present study was designed to tease apart these two strategies by providing conflicting feedback. Should feedback matter?

Domain General Problem Solving Strategies and the Differentiation of Means and Ends

Piaget (1954) suggests that infants in sub-stage 3 do not differentiate their actions, the means, or the ends. Because sub-stage 3 infants can focus upon their actions, we might expect them to be able to use trial-and error. Sub-stage 4 infants distinguish the ends but still see their actions and the means as undifferentiated. That is, they are not considering the relationship between the means and ends as causal. Rather, they focus on their own relationship to the ends. From this observation, we might expect sub-stage 4 infants to be able to use hill-climbing. By sub-stage 5, infants are able to distinguish their actions, the means, and the ends. They can recognize the special causal relationship between the means and ends as distinct from themselves. This suggests that sub-stage 5 infants can use means-ends analysis. From a Piagetian perspective, we might expect infants to gradually acquire domain general problem solving strategies in the following order: trial-and-error, hill-climbing, and means-ends analysis.

This Piagetian ordering of domain-general problem solving strategies is also sensible within an information processing framework. Consider the relationship between trial-and-error and hill-climbing. In order to use hill-climbing, we must recognize changes in the distance between our current-state and our goal-state. The most rudimentary way of recognizing the distinction is seeing the use of an operator as getting our current-state and goal-state to match. In this sense, we might expect trial-and-error to be a step on the way to using hill-climbing. Consider the relationship between hill-climbing and means-ends analysis. In order to use means-ends analysis, we must recognize a causal relationship. The most rudimentary way to begin to recognize a causal relationship is to first recognize a simple relationship between the operators and goal-state. A metric between the current state and the goal state is a salient relationship to focus upon. In this sense, we might expect hill-climbing to be a step on the way to using means-ends analysis. The present experiment is designed to test if domain-general problem solving strategies develop gradually or if they emerge all at once in a complete repertoire.

Hypotheses

The present study is designed to tease apart the use of hill-climbing and means-ends analysis strategies so that we can see if they develop. The present study also examines whether infants learn to use a strategy over the course of the experiment. The age-groups for the present study were chosen in accordance with Piaget's sub-stages in order to see how nearly infant performance corresponds to these stages. From these considerations, I suggest three hypotheses to be tested in the present study.

First, I predict that domain general strategies emerge gradually during infancy in the following order: trial-and-error, hill-climbing, and means-ends analysis. Second, I predict that these shifts are related to age. In particular, the transitions from Piaget's substages 3, 4, and 5 parallel the shifts between trial-and-error, hill-climbing, and means-ends analysis. Third, I predict that a basic mechanism for this change is learning through experience with related problems. That is, once infants use trial-and-error, they can learn to use hill-climbing and once infants use hill-climbing, they can learn to use means-ends analysis. These hypotheses are summarized in Table 1.4.

Table 1.4: Hypothetical Development of Domain General Problem Solving Strategies

Age of Infants	**Strategies in the Repertoire**	**Strategies Infants can Learn**
Sub-stage 3	Trial-and-error	Hill-climbing
Sub-stage 4	Hill-climbing & Trial & Error	Means-Ends Analysis
Sub-stage 5	Means-Ends Analysis, Hill climbing & Trial & Error	

To represent these hypotheses succinctly, and in operational terms, consider the table (Table 1.5) that compares infants' performance on a hill-climbing and a parallel means-ends analysis task. An infant is classified in one of three ways: unable to use the strategy, able to learn the strategy, able to use the strategy. I hypothesize that infants would be classified into 4 of the 9 cells. The cells can be explained developmentally from the upper-left cell, downward and across, to the lower right cell. Substage 3 infants are unable to use hill-climbing or means-ends analysis, but some of them are able to learn to learn hill-climbing. By substage 4, infants can use hill-climbing and, are now able to learn means-ends analysis. Finally, substage 5 infants are able to use both hill-climbing and means-ends analysis.

Table 1.5: Hypothetical Results of Task to Tease apart Hill-climbing & Means-Ends Analysis

Hypothetical Results		**Means-Ends Analysis Task**		
		Unable to Use	**Able to Learn**	**Able to Use**
Hill-Climbing Task	**Unable to Use**	Some infants, especially in sub-stage 3	No Infants	No Infants
	Able to Learn	Some infants, especially in sub-stage 3	No Infants	No Infants
	Able to Use	No Infant	Many infants, especially in sub-stage 4	Many infants, especially in sub-stage 5

Table 1.5 shows the basic structure for an experiment to test these hypotheses. I need two tasks that are as conceptually similar as possible, yet require different strategies to solve successfully. One is solvable by hill-climbing and one solvable only by means-ends analysis. The following method describes a task that meets these criteria.

Method

Participants

Infants were recruited from several middle-class towns in central Pennsylvania. Letters were sent to parents who had announced their child's birth in a major local newspaper. The 48 infants represent three age groups equally: 9-month-olds (M=9.2 SD=.3), 12-month (M=12.4 SD=.3) olds, and 16-month-olds (M=16.5 SD=.6). Equal numbers of male and female infants participated in each age group. Approximately 95% of the infants were Caucasian, the remaining 5% were from ethnic minority groups. Twelve other infants were excluded from the experiment: three because they were extra participants, five because of technical errors, and four because they failed to reach criteria before becoming fussy. No compensation for participation was offered but infants and their parents received a "Penn State Infant Problem-Solver Degree" with their infant's photograph and name as a thank-you gift.

Materials

Video cameras recorded each session. Parents held their infant in their laps while seated on a swivel chair. The swivel chair was positioned at the mid-point along the canal apparatus. The canal apparatus was 121 cm wide and 135 cm long. The canals were 8.5 cm wide by 6 cm high and the two experimental canals were separated by 16.5 cm. The wooden canal apparatus is covered with contact paper and regions of the canal within the infant's reach were also covered with foam. See Figure 2.1 for the arrangement of the canals. Obtuse, Right, and Acute canals are named for the type of angle formed at their intersections. The right canal is not used in this study; it was constructed on the apparatus in anticipation of future research. See Appendix I for a description of the complete geometric construction of the canals.

Cloths of a contrasting color were sewn to match the length and width of the canals. At one end of each cloth, small walls were sewn to the height of the canals. This was done to minimize chances of the toy falling off the cloth. If the toy became snared or fell off the cloth, the experimenter immediately put the toy back on the cloth. Toys were chosen by four criteria. They had to be: (1) age appropriate, (2) large enough to be clearly seen by the infant, (3) small enough to pulled without obstruction through the canals, and (4) capable of producing a sound that could attract the infant's attention, but not make sound while being pulled through the canal. The experimenter listened to a metronome through an earpiece to count out seconds.

Figure 2.1: Arrangement of the Canals on the Canal Apparatus.

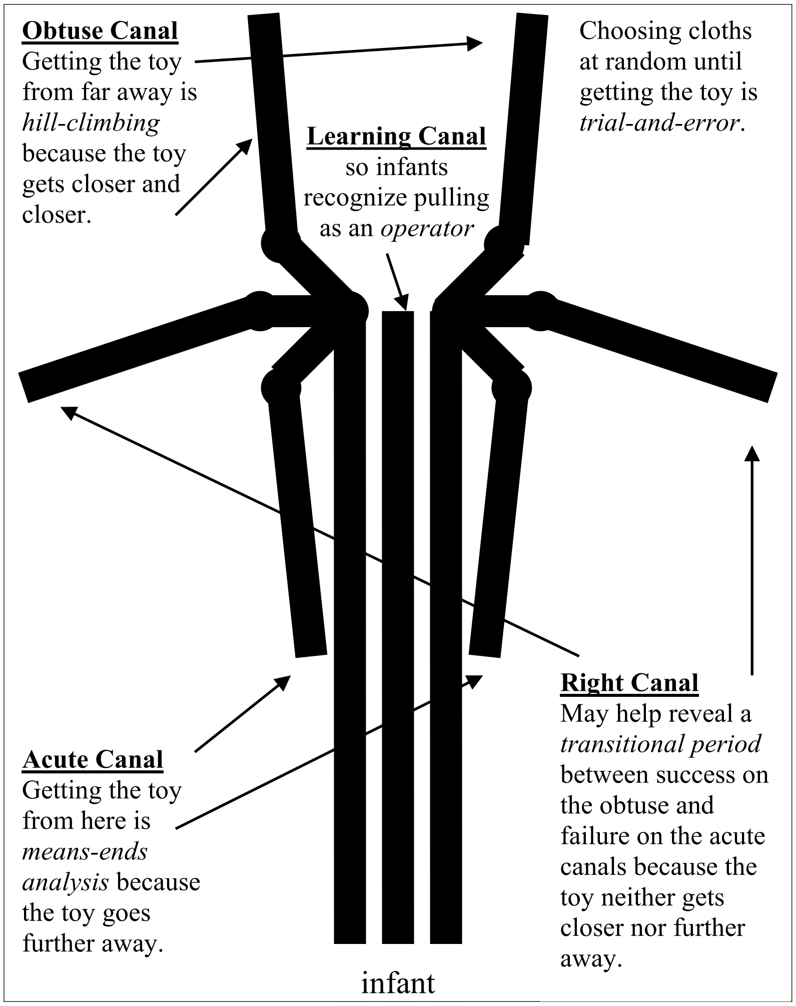

Design

Infants were tested individually in a within-subject experimental design. Each infant and parent dyad participated in a single forty-minute session. The experimental procedure took about

20 minutes. The remaining time included meeting the parent, discussing the experiment, teaching parents their role, and debriefing the parent after the experiment. The experiment consisted of two parts: the learning phase and the problem solving phase. The problem solving phase included two types of trials: acute angle trials and obtuse angle trials. Infants were block-randomized by age into a counterbalancing of trial order in the problem solving phase. See Appendix II for trial orders and a description of the constraints used to generate the psuedo-random counterbalancing orders.

For each trial within both phases of the experiment the infant-parent dyad sat on a swivel chair at one side of the canal apparatus and the experimenter stood at the opposite side of the canal apparatus. For each trial, the experimenter set down the cloth(s) and chose a toy at random. The cloths were laid in the canals reaching from the infant to the other end of the canal. At the experimenter's request, the parent swiveled around in the chair so that the infant was in front of the canals. The experimenter got the infant's attention by moving the toy and producing the toy's sound. The parent tossed the toy from the previous trial under the table to the experimenter. Once the infant's attention was fixed on the new toy, the experimenter placed the toy on the end of a cloth out of the infant's reach. For trials with two cloths (the problem solving phase), the experimenter moved both of her arms in symmetry as though placing the toy on both cloths. The experimenter did not look at her hands as she placed the toy. The parent and experimenter encouraged the infant verbally but did not gesture towards the toy. If the infant did not pull a cloth for 8 seconds, the experimenter asked, "I wonder if mom/dad can get the toy?" The parent then retrieved the toy by pulling the cloth in clear view of the infant. Upon retrieving the toy, the parent swiveled the chair so that the infant could not see the canal apparatus as the experimenter set up the next trial. This procedure is repeated for all trials in both the learning phase and the problem solving phase.

Learning Phase

Only one canal, the learning canal (see Figure 2.1), was used during the learning phase. The problem-solving phase began once the infant pulled the cloth two times within three trials. The learning phase insured that all infants could use the operator "pull the cloth", perceive distance, and perceive changes in distance. However, the infants did not necessarily understand how to choose among multiple cloths. They also did not necessarily understand the canal's role of guiding the cloth because the learning canal was straight.

Problem Solving Phase

On obtuse angle trials, two cloths were laid from the infant to the ends of both obtuse angle canals. A toy was placed on the left side for half of the obtuse trials and on the right side for the other half of the obtuse trials. On the acute angle trials, two cloths were laid from the infant to the ends of both acute angle canals. A toy was placed on the left side for half of the acute trials and on the right side for the other half of the acute trials. In total, there were 16 trials in the problem solving phase. Figure 2.2 illustrates the way cloths were laid in the canals for each trial type. Appendix II describes how trial order was counterbalanced.

Figure 2.2: Arrangement of Cloths in the Canals for Each Trial. Shaded regions represent the cloth. View is from above.

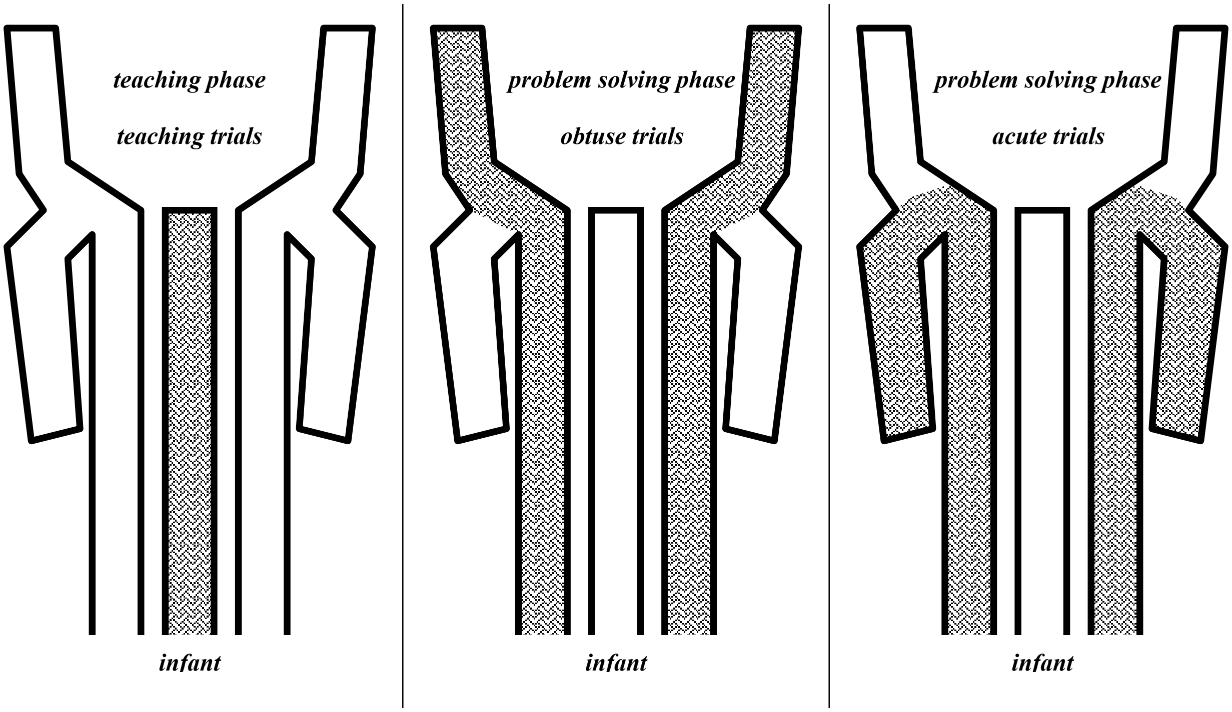

Measures

Each trial was coded from videotape by two raters. Raters classified each trial with a "yes" or "no" response to the following question: "Did the infant retrieve the toy?" Reliability was .86 as calculated using Cohen's Kappa. To avoid artificially inflating Kappa, this analysis did not include trials that were not completed by infant. That is, if an infant only completed 10 of the 16 trials, the remaining 6 empty cells were treated as missing data instead of coder agreement.

Results

No predictions have been made about counterbalancing condition, sex, and race. Furthermore, t-tests of these variables to predict total trials completed, total retrievals, and total pulls revealed no significant effects. These variables are not considered further. The results of an Analysis of Variance revealed no significant effect of age group predicting the total number of trials completed (F(2,45)=.243 p=.645). The subsequent results are divided into three parts. First, I examine the relationship between pulling the cloths and retrieving a toy, to determine if significant differences are likely to be due to executive function factors like persistence. Second, I examine how age and canal relate to retrievals for those infants showing any ability to retrieve the toy during the session. Finally, I examine the cloth pulling behaviors of those infants who never retrieved a toy during the session.

Relationship Between Pulling Cloth & Retrieving the Toy

To examine the relationship between pulls of the cloth and retrievals of the toy, I conducted a repeated measure ANOVA to predict total number of successful trials with Age Group (9mo, 12mo, 16mo) as a between subject factor and Canal Angle (obtuse, acute) and Successful Performance Type (pull, retrieval) as within subject factors. As expected, these was a main effect for Successful Performance Type because, by definition, it is impossible to retrieve the toy without pulling the correct cloth. Also as expected, there was a significant main effect of Age Group. This will be discussed later in the contexts of canal angle and half of the experimental session. Most notably, there was no interaction between age group and performance type (figure 3.1).

Figure 3.1 Successful Performance Type by Age Group. All 48 infants included: 16 9-month olds, 16 12-month-olds, and 16 16-month olds.

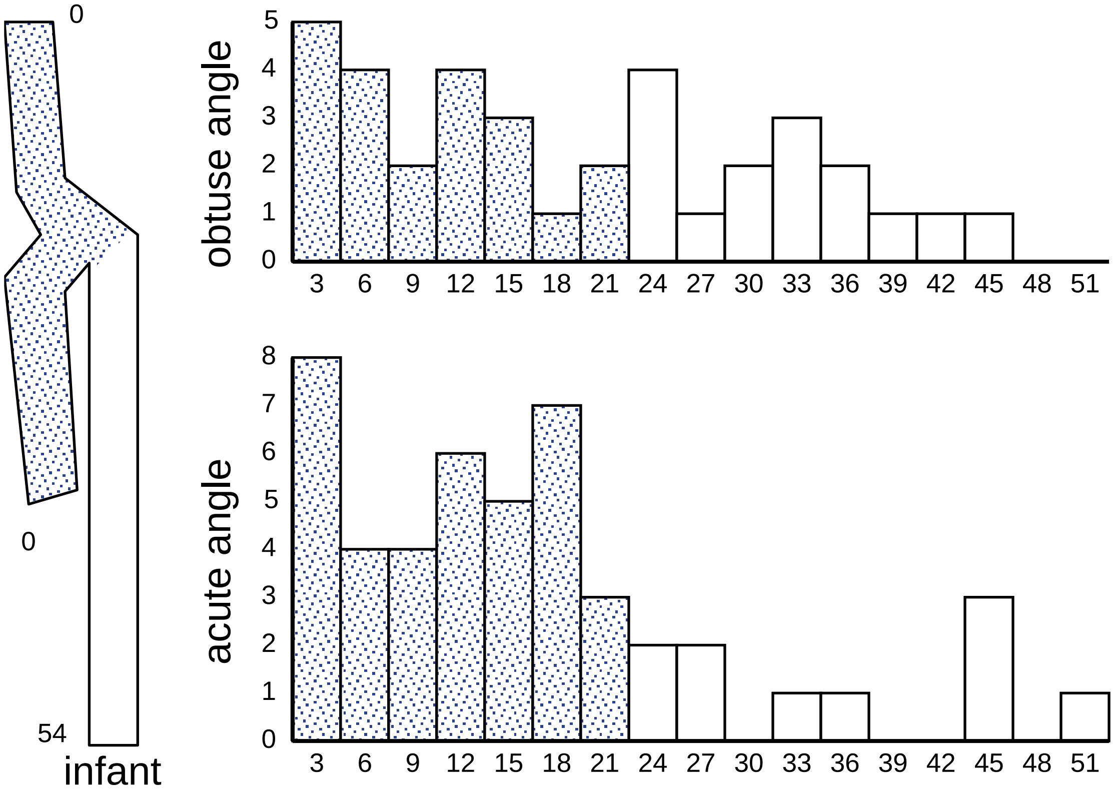

That is, once infants pulled the correct cloth, the chances that they would retrieve the toy were not related to their age groups. This suggests that any age-related differences were not due to persistence in pulling the cloth.

Age Group, Canal Angle, and Experimental Session Half

The gradual acquisition hypothesis makes predictions about the first half of the experimental session. The learning hypothesis makes predictions about the relationship between the first and second half of the experimental session. Before proceeding to these analyses, I examine the overall effects for canal throughout the experimental session (Figure 3.2). I conducted a repeated measure ANOVA to predict total successful retrievals with Age Group (9mo, 12mo, 16mo) as a between subject factor and Canal Angle (obtuse, acute) as a within subject factor. As expected, there were no significant effects.

Figure 3.2: Total Number of Retrievals across the experimental session as a function of age group and canal angle. 27 infants included. 21 infants excluded for failure to complete all 16 trials of the experimental session. Figure includes: 6 9-month olds, 10 12-month-olds, and 11 16-month olds.

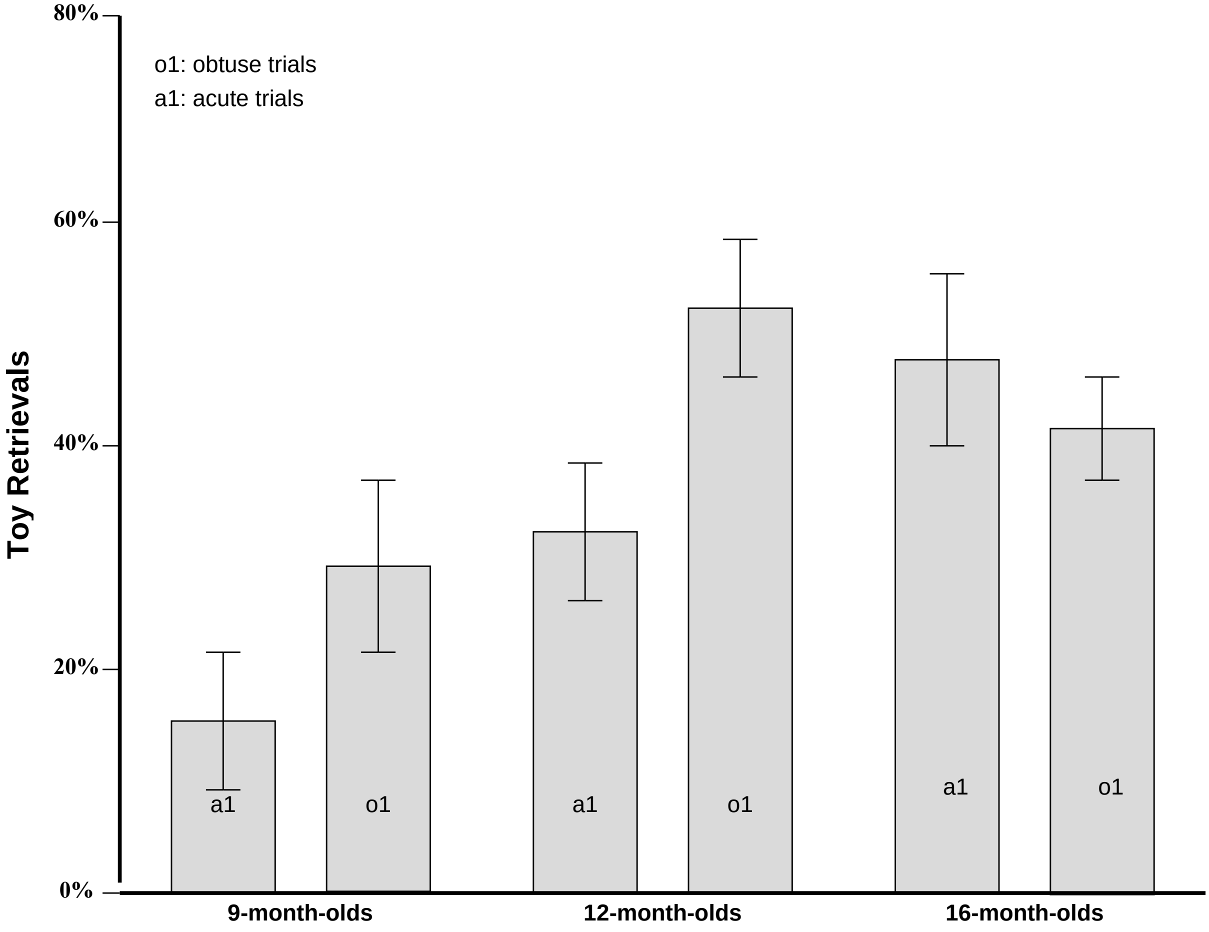

To examine the gradual acquisition hypothesis, I conducted a repeated measures ANOVA predicting number of retrievals with Age Group (9mo, 12mo, 16mo) as a between subject factor and Canal Angle (obtuse, acute) as a within subject factor. Tests of significance are treated as two-tailed unless: (1) there was an obvious direction of an effect expected from the analysis and (2) a significant effect in the opposite direction would falsify the hypothesis. To avoid possibly conflating the development of the ability to complete the task and the ability to solve the problems, only those infants who retrieved the toy at least once were included in this analysis. The nine infants who never retrieved a toy will be examined later. Results reveal a main effect of Canal Angle F(1,38)=8.525 p=.003 (one-tailed), and a marginally significant Age Group by Canal Angle Interaction F(2,38)=2.277 p=.058 (one-tailed).

Figure 3.3: Total Number of Retrievals across the first half of the experimental session as a function of age group and canal angle. 41 infants included. 7 infants excluded for failure to retrieve a toy on any trials of the experimental session. Figure includes: 10 9-month olds, 16 12-month-olds, and 15 16-month olds. Graph including only those infants who completed all 16 trials is comparable.

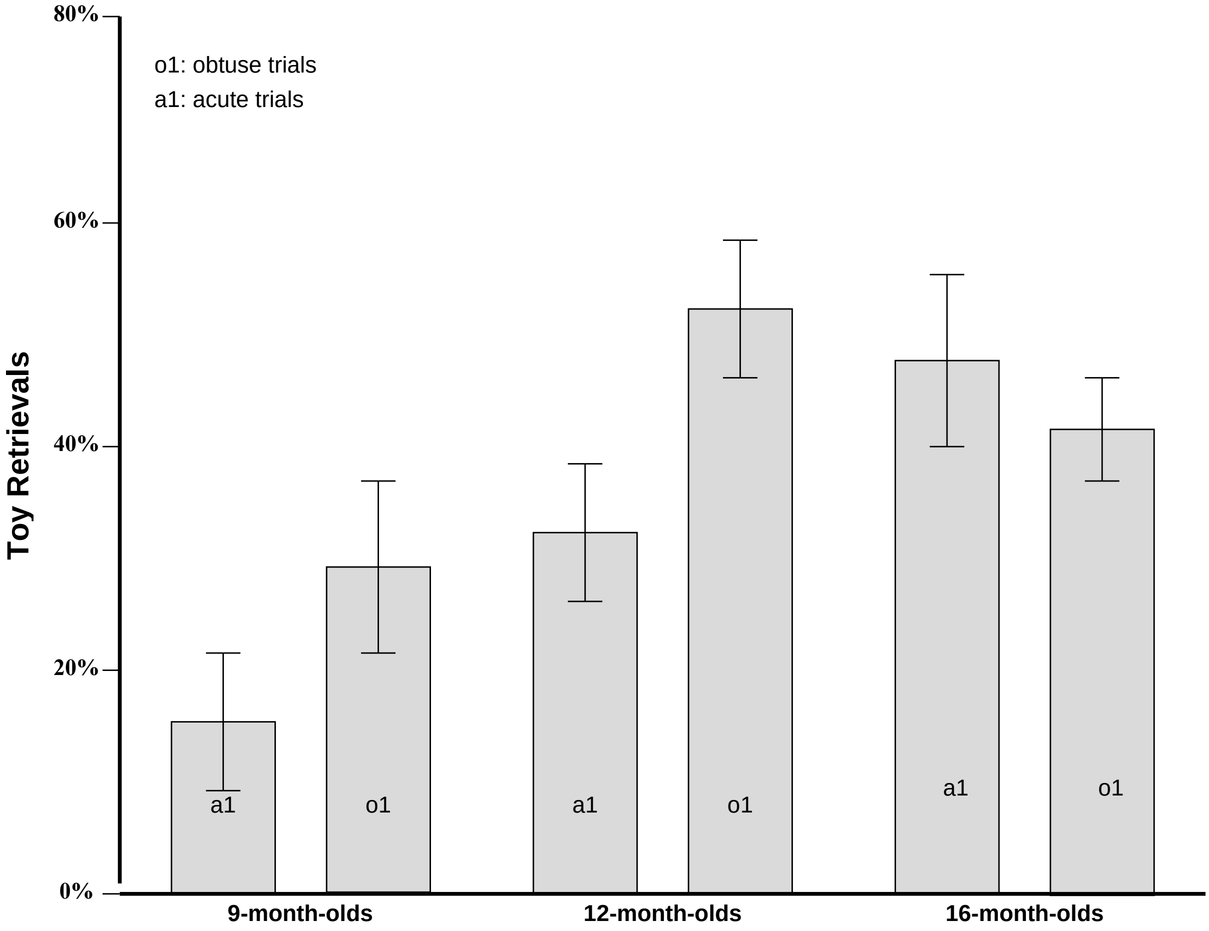

To test the learning hypothesis, I conducted a second ANOVA with all those infants in the preceding ANOVA who completed the entire experimental session. This allows a comparison of the halves of the experimental session (first, second) as a within-subject factor and Age Group (9mo, 12mo, 16mo) and Canal Angle (obtuse, acute) as between-subject factors. The main effect for Canal Angle becomes marginal F(1,24)=1.106 p=.075 (one-tailed). The marginally significant Age Group by Canal Angle Interaction F(2,24)=2.451 p=.053 (one-tailed) remained. The mean number of retrievals for each age group and canal angle also remained the same (figure 3.3).

There was a significant main effect for Experimental Session Half F(1,24)=9.524 p=.002 (one-tailed) and a marginally significant Canal Angle by Experimental Session Half interaction F(1,24)=1.769 p=.098 (one-tailed). That is, the number of retrievals increased from the first to the second half of the session and this increase was more pronounced for the acute canal angle trials than for the obtuse canal angle trials (Figure 3.4)

Infants who never Retrieved a Toy

Nine of the 48 infants never retrieved a toy during the experimental session. I examined whether infants who never retrieved the toy were nevertheless increasingly pulling the cloth and whether infants pulled cloths more on obtuse canal trials than on acute canal trials. Because the sample size is small, it is difficult to conduct a statistical test with multiple factors. I conducted a paired-sample T-tests of each half of the session, both overall and separated by canal angle. There was neither a significant difference in pulls between the first and second half of the session, nor a significant difference between the canal types within the first half of the session. However, there was a marginally significant difference in pulls for the second half of the session $F(1,4)=1.633$ $p=.089$ (Figure 3.5).

Figure 3.5 Total number of pulls of the toy cloths comparing halves of the experimental session as a function of canal angle. 5 infants included of the 7 infants who never retrieved the toy. 2 infants are excluded for failure to complete all 16 trials. Figure includes: 5 9-month olds, 1 12-month-olds, and 1 16-month olds.

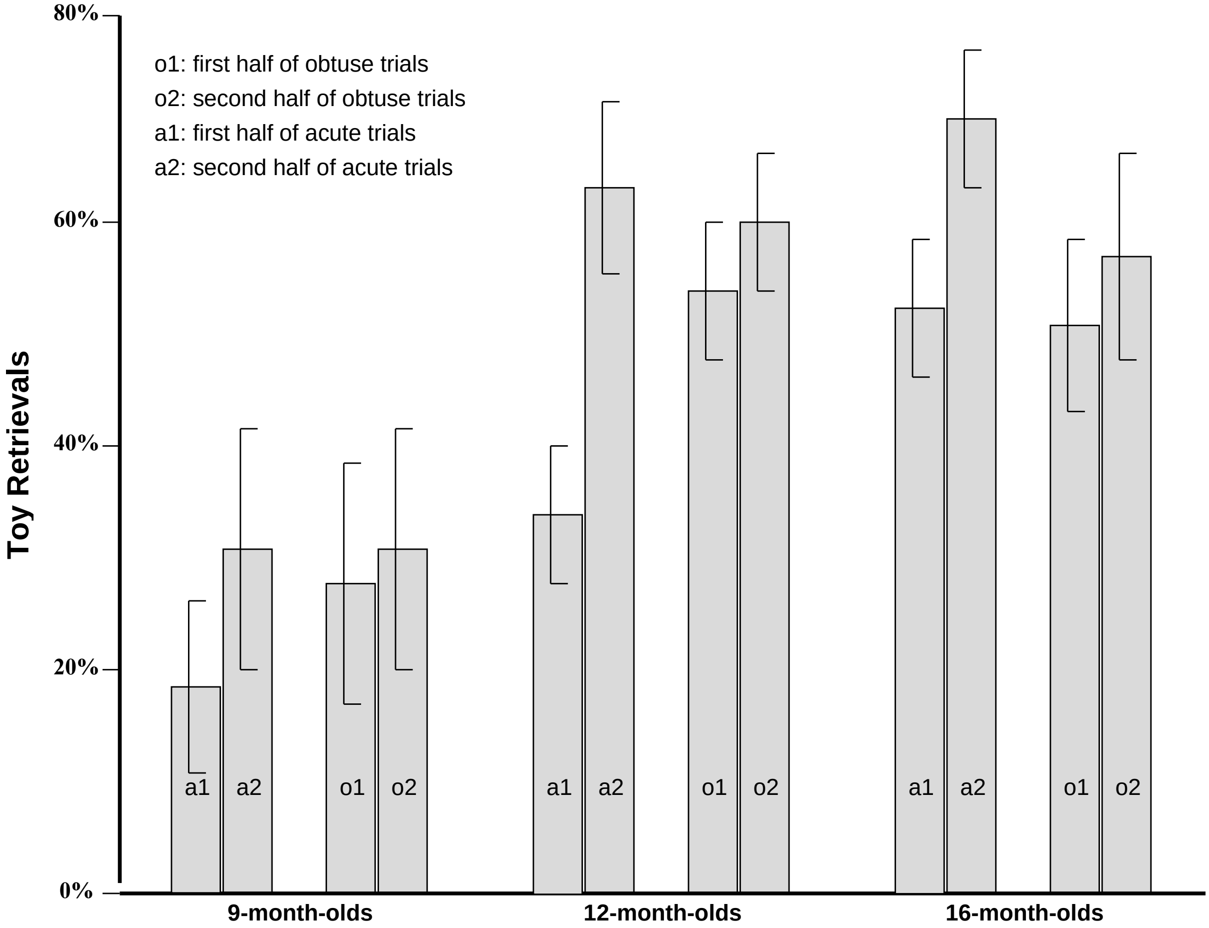

Discussion

The goal of the present study was to examine the origins of hill-climbing and means-ends analysis during infancy. The present study represents the first time that infant's use of means-ends analysis have been studied directly, with conceptually identical tasks that differ in the domain general problem solving strategies that can be used to solve the problem. In addition to the direct assessment of strategy use, this study unambiguously measured success retrieving the toy. The results showed that 9 months olds did not typically use hill-climbing or means-ends analysis in the task. Sixteen month olds used both strategies. The 12 month old infants used hill-climbing just as frequently as the 16 month old infants did, but they were no more likely to use means-ends analysis than the 9 month old infants were. This suggests that the ability to use both types of domain general problem solving strategies develops gradually throughout infancy. Furthermore, closer examination of the 12 month old infants showed that following 8 trials on the problem requiring means-ends analysis, they were able to solve the problem as well as 16-month-olds. This suggests that some infants can learn means-ends analysis within a particular context. More broadly, these results support Piaget's observation that infants gradually differentiate the means from the ends. But, the results are inconsistent with Willatt's claim that 9-month-olds can use means-ends analysis, and with the implicit assumption about infant's problem solving in Case's Neo-Piagetian model.

Recall from the introduction that it is possible to illustrate results consistent with the hypotheses in a succinct table (Figure 1.6 & Figure 4.1). The results of the present study could be captured in this table, although somewhat arbitrary categorical distinctions in performance need to be made. For purposes of illustration, I fit the current study into this table to help explain the meaning of the statistical results (Figure 4.2).

Table 4.1: Hypothetical Results of Task to Tease apart Hill-climbing & Means-Ends Analysis

Hypothetical Results		**Means-Ends Analysis Task**		
		Unable to Use	**Able to Learn**	**Able to Use**
Hill-Climbing Task	**Unable to Use**	Some infants, especially in sub-stage 3	No Infants	No Infants
	Able to Learn	Some infants, especially in sub-stage 3	No Infants	No Infants
	Able to Use	No Infants	Many infants, especially in sub-stage 4	Many infants, especially in sub-stage 5

Table 4.2: Actual Results of the Present Experiment

Descriptive Breakdown of Infant Performance		**Acute Canal**		
		Can't do	**Learns**	**Can do**
Obtuse Canal	**Can't do**	9 infants (19%) 7 nine mo 1 twelve mo 1 sixteen mo	2 infants (4%) 0 nine mo 0 twelve mo 2 sixteen mo	1 infant (2%) 1 nine mo 0 twelve mo 0 sixteen mo
	Learns	2 infants (4%) 1 nine mo 1 twelve mo 0 sixteen mo	6 infants (13%) 1 nine mo 1 twelve mo 4 sixteen mo	1 infants (2%) 0 nine mo 1 twelve mo 0 sixteen mo
	Can do	1 infant (2%) 1 nine mo 0 twelve mo 0 sixteen mo	8 infants (17%) 2 nine mo 5 twelve mo 1 sixteen mo	18 infants (38 %) 3 nine mo 7 twelve mo 8 sixteen mo

To do this I captured each infant's performance during the study using least-square best-fit lines of trial number (0 to 7) to predict retrieval (1 for yes, 0 for no). The y-intercept values provided a starting ability and the slope provided a learning rate. If a starting point was greater than or equal to 1/3 then the infant was classified as able to use the strategy. Of those infants not able to use the strategy (i.e. the infants not yet classified), those with a positive slope were classified as able to learn the strategy. Those infants who started the study with retrievals below 1/3 and did not have a positive slope were classified as unable to use the strategy.

Table 4.1 shows that we would expect all infants to fall into 4 of the 9 possible cells. In actuality (Table 4.2), 79% of the infants fall into these cells. This is because, unexpectedly, there are 6 infants who learned both strategies (middle center cell) instead of being unable to do the acute trials while able to learn obtuse trials (middle left).

Nineteen percent of the infants, mostly comprised of 9 month olds, fell into the upper-left cell. These infants were unable to retrieve the toy on both obtuse and acute trials.

Seventeen percent of the infants, comprised mostly of 12 month olds, fell into the lower center cell. These infants were able to use hill-climbing and able to learn to use means ends analysis. As previously discussed about 12 month-old infants in general, this supports the learning hypothesis for means-ends analysis. As shown earlier (Figure 3.3), nine and twelve-month olds began the study able to use hill-climbing but not means ends analysis. Half of the sixteen month olds began the study able to use both hill-climbing and means-ends analysis. In the course of the study the 12 month old infants were able to learn to use means-ends analysis.

These results also support the gradual acquisition hypothesis. Nine-month olds had difficulty using both strategies but 12-month-olds were able to use hill-climbing without being able to use means-ends analysis. Finally, by 16-months-old, infants were able to use both strategies. This is consistent with the 38% of infants, the majority of whom are 16 months old, who were classified as falling into the lower-right cell. Results of this study are consistent with the notion that domain general problem solving strategies develop gradually as part of a global developmental process. Nevertheless, a single experimental task does not demonstrate this possibility. Future research could relate how infants acquire domain-general problem solving strategies with how infants and toddlers develop an increasing flexibility to use domain-general strategies by using analogical-transfer. Infants could participate in two different tasks that disentangle the domain-general strategies as in the task in the present study.

The hypotheses of this study were not entirely supported by the results. On average, nine month old infants failed to learn to use hill-climbing during the experiment. Perhaps the task and environment, including an especially long cloth winding through an elaborate table, an unfamiliar experimenter showing toys but refusing to provide them directly, and a parent encouraging them to try and get the toy, was simply too overwhelming for such young infants to learn. This lack of support for the hypothesis parallels the unexpectedly small number of infants in the middle-left cell of Table 4.2. Few of the infants who were learning hill-climbing remained incapable of learning means-ends analysis. Instead, 13% of all infants were able to learn both strategies simultaneously. Perhaps this is the result of conditioned learning. That is, infants may simply have formed an association between pulling the cloth (stimulus) and getting toys (response).

But, if infants are able to learn to use both strategies *simultaneously*, how could this study have shown that these strategies emerge at *different* times? The key to resolving this paradox may be to note the difference between the experimental session and infant's typical life experiences. In the experiment, infants learned with conceptually very *similar* problems. In their typical problem-solving experiences, infants must learn to use problem-solving strategies as a *generalization* of many distinct types of problems. Both strategies could be learned at the same time in this experiment because the exact same task was used repeatedly. In this sense, it remains unclear whether these infants were really learning to use the domain general problem solving strategy or if they were learning a domain-specific strategy for cloths-through-canal problems.

Data from those infants who never retrieved a toy can help illuminate this issue because they never had the opportunity to learn the particular stimulus-response relationship: pulling cloth gets the toy. As described previously (Figure 3.4), even without getting the toy, these infants still pulled on the cloth. Notably, as the session proceeded they became more likely to pull on the obtuse trials but they did not increase the amount that they pull on the acute trials. Perhaps these infants are beginning to learn to use hill-climbing since they increase their pulls despite not getting the toy when a pull moved the toy closer to them. Put within behaviorist concepts, perhaps difference-reduction is rewarding just like achieving a goal is rewarding.

Origins of Hill-Climbing

The pattern of results in this study suggests that hill-climbing may be an innate or early emerging ability that becomes available to infants as soon as they conceptually understand a problem solving context. Is there other evidence suggesting hill-climbing may be innate? How might the innateness of hill-climbing ability in infants be tested directly?

Figure 4.3 Chapuis (1987) Tasks given to horses, cats, and dogs.

1. less divergent & shorter way	2. less divergent but equal length	3.equal divergence and shorter length	4. more divergence and shorter length
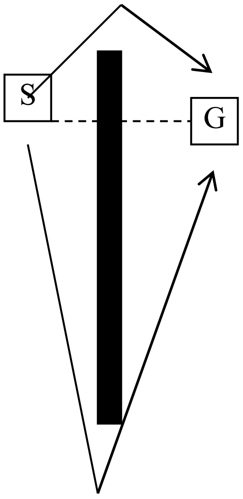	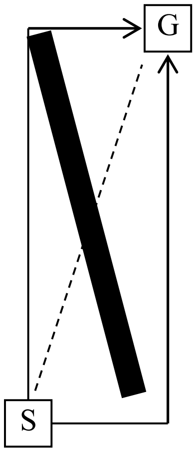	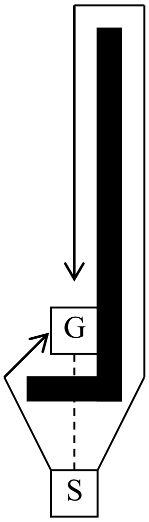	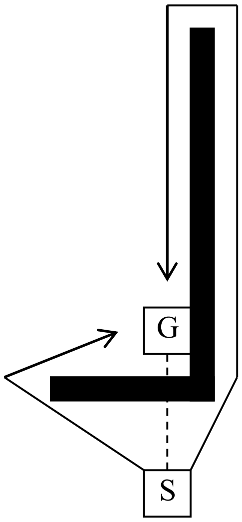
Unlike the 'incorrect' path, the 'correct' path always moves closer to the goal.	*Tthe 'correct' path moves closer to the goal & then away. The incorrect path is the reverse.*	*Unlike the 'incorrect' path, the 'correct' path always moves closer to the goal.*	*Both the 'incorrect' and 'correct' pathes move further from the goal to eventually reach it*

There are several examples of animal research that may suggest hill-climbing is an innate ability of many animals. For example, Gallistel (1990) had mice try to find cheese in a maze. The mice were more likely to make a wrong turn if it was a wrong turn moving closer to the cheese than a wrong turn moving away from the cheese. This suggests mice may be using hill-climbing. However, this task does not tease apart a mouse's use of hill-climbing from means-ends analysis. The animal literature has not traditionally used the language of information processing so, at least to my knowledge; this has never been directly tested with animals. However, some research on “divergence” may be particularly relevant. “Divergence” is the angle a participant needs to deviate from a direct path to a goal in order to eventually reach that goal. Chapuis (1987) conducted a particular well designed study to tease animals’ abilities apart to use divergence and distance information when pursuing a goal. She placed an obstacle between an animal and its food. The animal had to choose one of two possible ways around the obstacle to the food. She found that dogs, cats, and horses would: (1) choose a path that is less divergent and shorter more often than chance, (2) would choose a path of less divergence but equal length more often than chance, and (3) would choose a shorter path when divergence in equal more often than chance. However, these animals performed at chance when one path was more divergent but short in length. Figure 4.3 shows these obstacle tasks and provides a means-ends interpretation of each task. Notice how the animals were always able to choose the most direct path *except* when the most direct path meant moving further from the goal in order to reach it. This may suggest that these animals were instinctively able to use hill-climbing but did not have the ability to use means-ends analysis. Similar results have been found with other animals such as hamsters (Buhot & Poucet, 1987).

There are several possible future directions to help determine if hill-climbing is an innate ability. First, we could measure the rate of conditioned learning when the response is either: (1) achieving a goal, (2) moving closer to a goal but not achieving it, and (3) moving further from a goal and not achieving it. If hill-climbing is innate, we could expect infants to learn in conditions 1 and 2 before they learn in condition 3. Another future direction would be to replicate Chapuis (1987) with infants. This could be particularly interesting because we could do a further comparison between human infants and other animals. Consistent with the results of the present study, I would expect human infants to be able to learn to solve Chapuis (1987) obstacle 4 because this would suggest learning means-ends analysis. If we also replicated the study with other animals, would they be able to learn means-ends analysis with experience in this context too?

Variation Among, Between, and Within Infants

Despite the significance of age group differences, there were also considerable differences within age-groups (Table 4.4). For example, though as a group 9 month olds were the least successful group at retrieving the toys and as a group 16 month olds were most successful, there were 9 month olds who successfully retrieved the toy on 15 out of 16 trials and there were 16 month olds who never retrieved a toy. What accounts for this wide variability in performance?

Figure 4.4 Total Retrievals by Age Group

Total Retrievals	9 month olds	12 month olds	16 month olds
0 to 3	9	2	3
4 to 7	4	3	3
8 to 11	1	8	6
12 to 16	2	3	4

One possibility may be examined by looking at the structure of the experiment. In every trial, each infant eventually received the toy. This was either because of the infant's problem solving efforts, or because the experimenter cued the parent to solve the problem. Parents only retrieved toys at the experimenter's cue, so infants' attempts to induce their parent to get them the toy would not be successful. However, if infants persisted at trying to get their parent's help for the full 8 second delay, then it could appear as though they were successful when the parent started pulling. Future analyses will include an examination of the role of help-seeking behaviors in problem solving strategy selection.

Desire for the Toy

In order to study the emergence of domain general problem solving strategies, I designed this experiment to induce a specific goal in each of the children: a desire to retrieve the toy. Despite this effort, infants may or may not have adopted that goal. This appears likely because even the 15 month old infants were not performing at ceiling. There were examples that demonstrated that infants sometimes had different goals. For example, one 12 month old girl retrieved the cloth without the toy and used it as a scarf. A 15 month old boy who very enthusiastically retrieved a toy car on one trial did not retrieve a new toy on the next trial. Instead, he looked behind his mom and under the table in an effort to find out what happened to the car. Future analyses will examine how an infant's goals on a particular trial relate to their abilities to use the problem solving strategies.

Help-Seeking and Help-Giving

Parents could not be completely uninvolved in the experimental procedure because they needed to swivel in their chair, play with their children between trials, and take the toys away from their children. Consequently, parents were another source of variability. Parents engaged in subtle and not so subtle behaviors to encourage their children. For example, though parents were instructed to center their child between the canals, they seemed to have a tendency to lean just a bit more towards the canal with the toy. This is not a confound because parents were surprised to learn the hypotheses after their session so it is unlikely that any variance due to parental behaviors systematically favored or negated the hypotheses. Parental reactions to their infant’s ‘off task’ behavior also varied. Parents said things like, “Can you get the toy?” One father, seemingly bothered by his daughter looking around the room knocked on the table to reorient her attention to the table and cloths. In future studies, I may control parents’ behaviors more strictly. This is difficult because children experience some frustration at not having a toy that they want. If their parent’s behavior is notably different from usual, this may be an added source of stress. Even without this added stress and with these subtle cues infants did not perform at ceiling.

Working Memory Capacity

One hypothesis of this experiment concerned a basic mechanism underlying development. There is evidence that infants' learning acquire means-ends analysis through experience in solving similar problems. One possible source of variation between infants is improvement in working memory capacity. This mechanism is plausible because cognitive load is greater for means-ends analysis than for hill-climbing (Sweller, 1988) and individual difference in working memory capacity are related to problem-solving task performance (e.g., Rabinowitz et al, 2002). Moreover neurological development during infancy is thought to lead to greater working memory capacity (e.g., Bauer & Hertsgaard, 1993), Future research might compare infant's working memory capacities to their performance on tasks like this one. I would hypothesis that increases in working memory ability would predict more advanced problem solving strategy ability. There are already ways of measuring working memory capacity during infancy (e.g., Gilmore & Johnson, 1995; Schwartz & Reznick, 1999). However, these ways of measuring working memory capacity require groups of infants. To my knowledge, there is no developed methodology for assessing a particular infant's working memory capacity at a particular age. This methodological problem needs to be overcome to pursue this possible future direction.

Development of Strategies in Stages

Because the results of the present study suggest that domain-general problem solving strategies are acquired gradually, we might envision a strict stage-like model of domain general problem solving strategy development. This might seem like a sharp contrast to some developmental models of domain specific problem solving strategies, such as Siegler's Overlapping Wave Model (Siegler & Shipley, 1995). Siegler’s model represents children as always having many strategies from which to choose. Children test all of these strategies as though they are competing with one another to be used. Over time, this process of competition works akin to a process of natural selection and thus different principle strategies emerge as children develop (Siegler, 1998).

There are at least two distinct ways of addressing the inconsistency between these stage-like and non-stage-like models of strategy acquisition. First, though it would be more parsimonious to find the same basic developmental course underlying all strategy acquisition, there is no logical necessity that the basic developmental course underlying domain-general problem solving should be the same as the basic developmental course underlying domain-specific problem solving. If research findings suggest these different developmental courses, future research questions might address a broader model to integrate the strategies into a single model. Perhaps a unified model could note similarities between domain-general strategies and some domain-specific strategies, especially those used when children lack prior knowledge of the problem situation. This possibility is promising because previous research on domain-specific problem solving suggests that children may appear more stage-like when they have little conceptual knowledge of a domain (Siegler, 1983).

Focusing on experimental methods is a second way of addressing the possible inconsistency. The stage-like results hypothesized in this thesis may be more of an artifact of the experimental design than a true suggestion that the process is stage-like. Success on a canal type problem indicates an ability to use at least a single strategy; it does not distinguish which strategy was really chosen by the infant. For example, though I interpret an infant's successful retrieval in the obtuse-canal as the ability to use hill-climbing, I do not know if that infant is always using hill-climbing on those trials or if that infant switches between hill-climbing and means-ends analysis randomly. A possible future direction is to design a problem solving task that can fully disentangle which of the three domain general strategies an infant is using on a single trial. This new task combined with a micro-genetic approach could help us understand if the developmental course of domain-general problem solving really is stage-like.

Potential Application of this Research

Clinicians have concerns about cognitive disabilities not being detected by traditional screening and assessment measures of infants (Caron, Caron, & Glass, 1983). Though early biological factors (e.g. prematurity, respiratory difficulties) predict later cognitive disabilities, most infants with these risk factors do not develop cognitive disabilities. Further, most children are not diagnosed with learning disabilities until they enter schools and most children with learning disabilities have normal infant histories (Wyly, 1987). The use of information processing procedures is a viable alternative to traditional developmental assessments (Benasich & Bejar, 1992; Rose et al. 1991). I believe this task designed for this experiment may be particularly useful as an early assessment tool of learning disabilities because the task demands are modeled after the problem solving situations typical of school-age children and adults. Future research might include a longitudinal study to determine if performance on this task predicts later learning disabilities and IQ.

Table 4.5: Summary of the Milestones & Mechanisms of Domain General Problem Solving Strategy Development

Age	Strategy	Explanation of Milestones
6 months	no strategy	Infants can learn trial-and-error with experience solving problems but they can not learn other strategies because they lack conceptual pre-requisites. This period has not yet been tested empirically.
9 months	trial-and-error	Infants can not learn means-ends analysis because they lack conceptual pre-requisites. This experiments suggests these infants can not learn hill-climbing either.
10 to 11 months		Either infants become capable of learning hill-climbing or they spontaneously become capable of using this strategy. This period has not yet been tested empirically.
12 months	hill-climbing	Infants can learn means-ends analysis with experience solving problems. They do not learn other strategies because they have previously acquired those strategies.
15 months	means-ends analysis	The infants have acquired these three domain general problem solving strategies for the problem solving situation.

This experiment supports the notion that domain general problem solving strategies develop gradually during infancy. Learning accounts for the development of means-ends analysis but not for the development of hill-climbing. Figure 4.5 revises the model of domain general problem solving development based on the results of the present experiment.

Problem solving during infancy has been observed in a wide range of contexts and it has been studied from a wide range of theoretical perspectives. Previous research on infant problem solving has been extensive albeit without problem solving typically as the primary research question. This thesis brings this diverse research together to begin to address the basic developmental question of when and how problem-solving strategies originate. Yet, the results of this study open up more questions than it solves. The future directions proposed may lead to increased knowledge about the stage-like nature of problem solving development, the possibilities of transferring domain-general strategies across contexts, and further the exploration of possible basic mechanisms underlying cognitive development.

Appendix A: Geometric Construction of the Canals

Figure A.1: Geometric Construction of the Canals

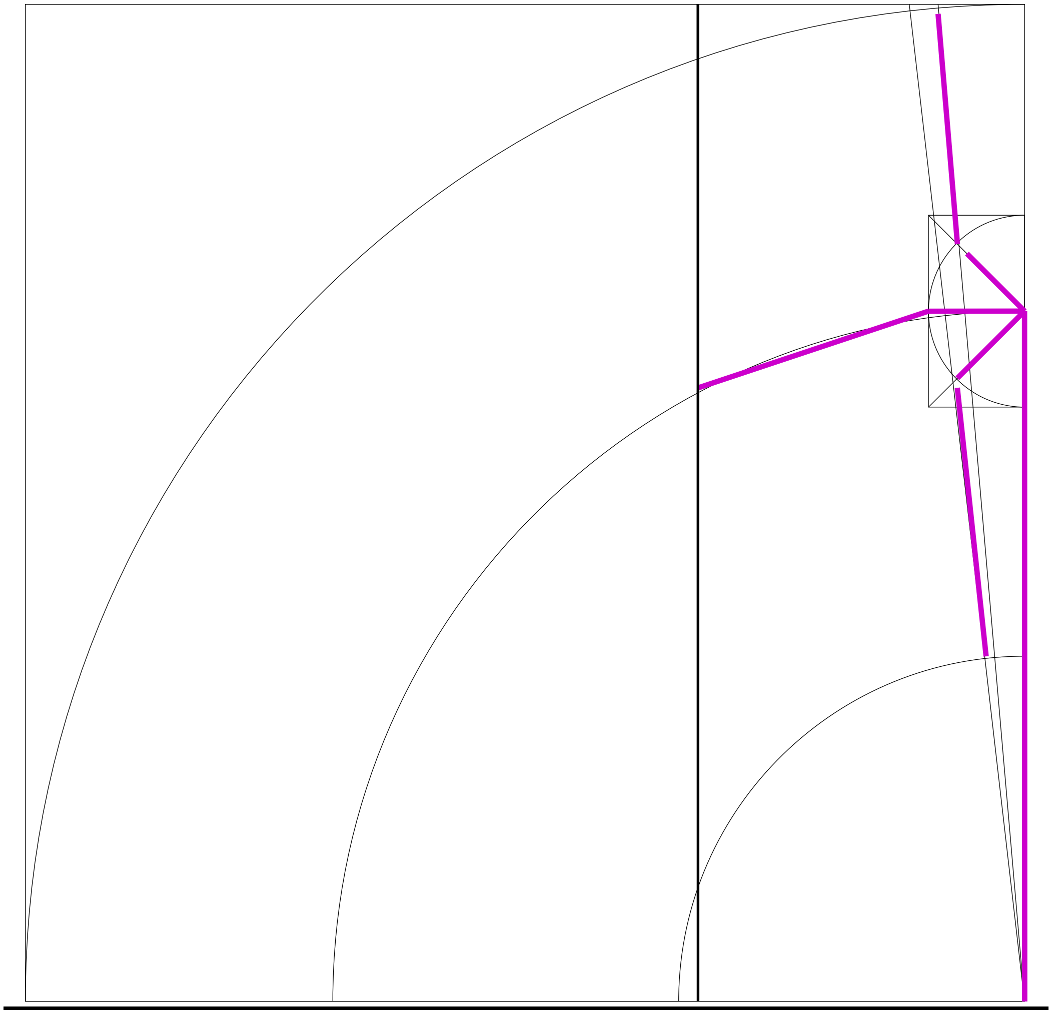

The width of the canals is not shown in the geometric construction to insure clarity. Each experiment canal consists of three paths: an outermost, middle, and innermost path. Each corresponding path for each canal is the same length. All three canals share an innermost path. The outermost paths of the obtuse and acute canals are constructed such that they fall along radii from the infant's center. The right angle canal is constructed to fit the circumference of a circle whose radius is the infant's center. The middle paths are just long enough to provide the necessary curvature. This may not be clear from the construction because width is not shown. The middle paths are radii of a circle whose center is the edge of the innermost path. The angles of between the middle and innermost paths for are 45° away from the infant for the obtuse canal, 45° towards the infant for the acute canal, and 90° for the right canal.

Appendix B: Counterbalancing

Figure B.1: The Canals with Abbreviations

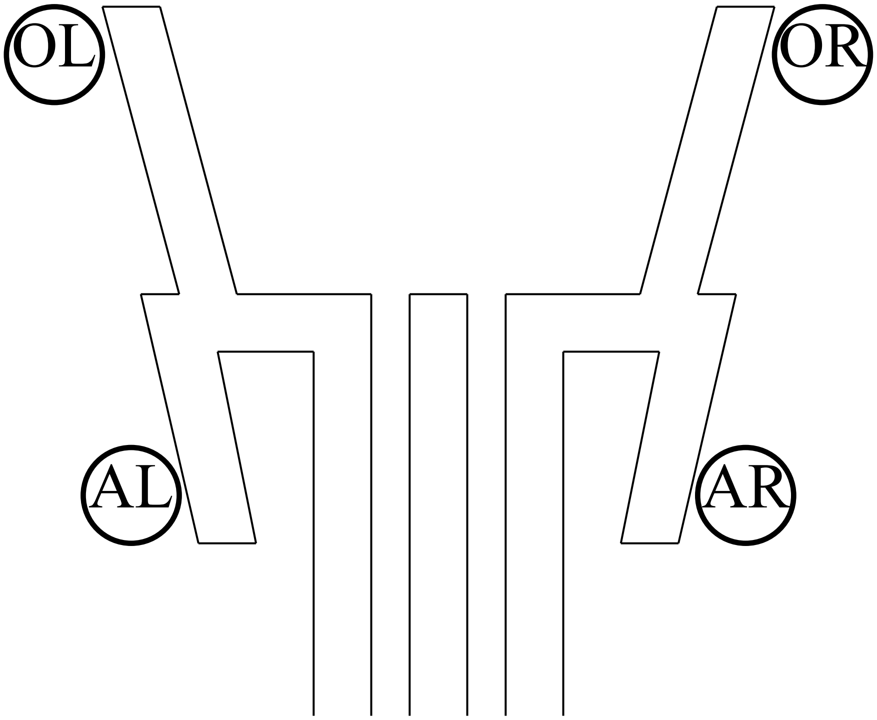

The Problem Solving Phase of the Experiment consists of 16 trials. These trials are divided into four blocks of four trials each. Within a block, there is one each of four types of trials: Obtuse Canals when Toy is Placed to Left (OL), Obtuse Canals when Toy is Placed to Right (OL), Acute Canals when Toy is Placed to Left (AL), Acute Canals when Toy is Placed to Right (AR).

The psuedo-random order of trials is subject to the following constraints: no more than two trials in a row are to the left, no more than two trials in a row are to the right, no more than two trials in a row are with the acute canals, no more than two trials in a row are with the obtuse canals, and no two trials in a row are the same position (exp: AR, AR).

The following are the actual counterbalancing orders used, F (forward) and R (reverse)

F: OL AR OR AL AR OL OR AL OR AL AR OL OR AL OL AR

R: AR OL AL OR OL AR AL OR AL OR OL AR AL OR AR OL

Appendix C: Pilot Study

A small pilot study was conducted before proposing this Master's thesis. This method of the pilot experiment closely resembles that of the proposed study. The key hypotheses are also the same: domain general problem solving strategies emerge gradually and domain general problem solving strategies are learned within a context. The pilot study was conducted to refine the method, and to insure that the hypotheses are plausible.

Method

Participants

Five infants, three male and two female, between 12 and 13 months of age (mean=391 days, s.d.=8 days) participated in the pilot study. The sample was chosen from a white rural middle-class sample. This age group was chosen because infants were expected to be capable of using hill-climbing and not capable of using mean-ends analysis.

Materials

The materials are identical to those in the proposed study. They include: video equipment, the canal apparatus, cloths, and toys.

Design

With three adjustments, the method of the pilot study is the same as the method of the proposed study. The first adjustment was the number of blocks. The pilot study included up to six blocks if the infants did not become disengaged sooner. Four blocks are included in the proposed study because participants in the pilot study persisted until after the fourth block and no infant's performance improved after the fourth block. The second adjustment is the counterbalancing that was somewhat different because of the different number of blocks. The third adjustment concerns how infants return the toy they retrieved on the preceding trial. In the pilot study, infants returned the toy to the experimenter before the parent swiveled his or her chair around. Losing the toy frustrated many infants and the results could be influenced by parents' differing skills in teaching their infants to regulate their emotions (e.g. Cole, Michel, & Teti, 1994). During the pilot study, the procedure was refined so that infants kept the toy as their parent around to face the canals.

Results

Figure C.1: Average Performance of Infants in the Pilot Study

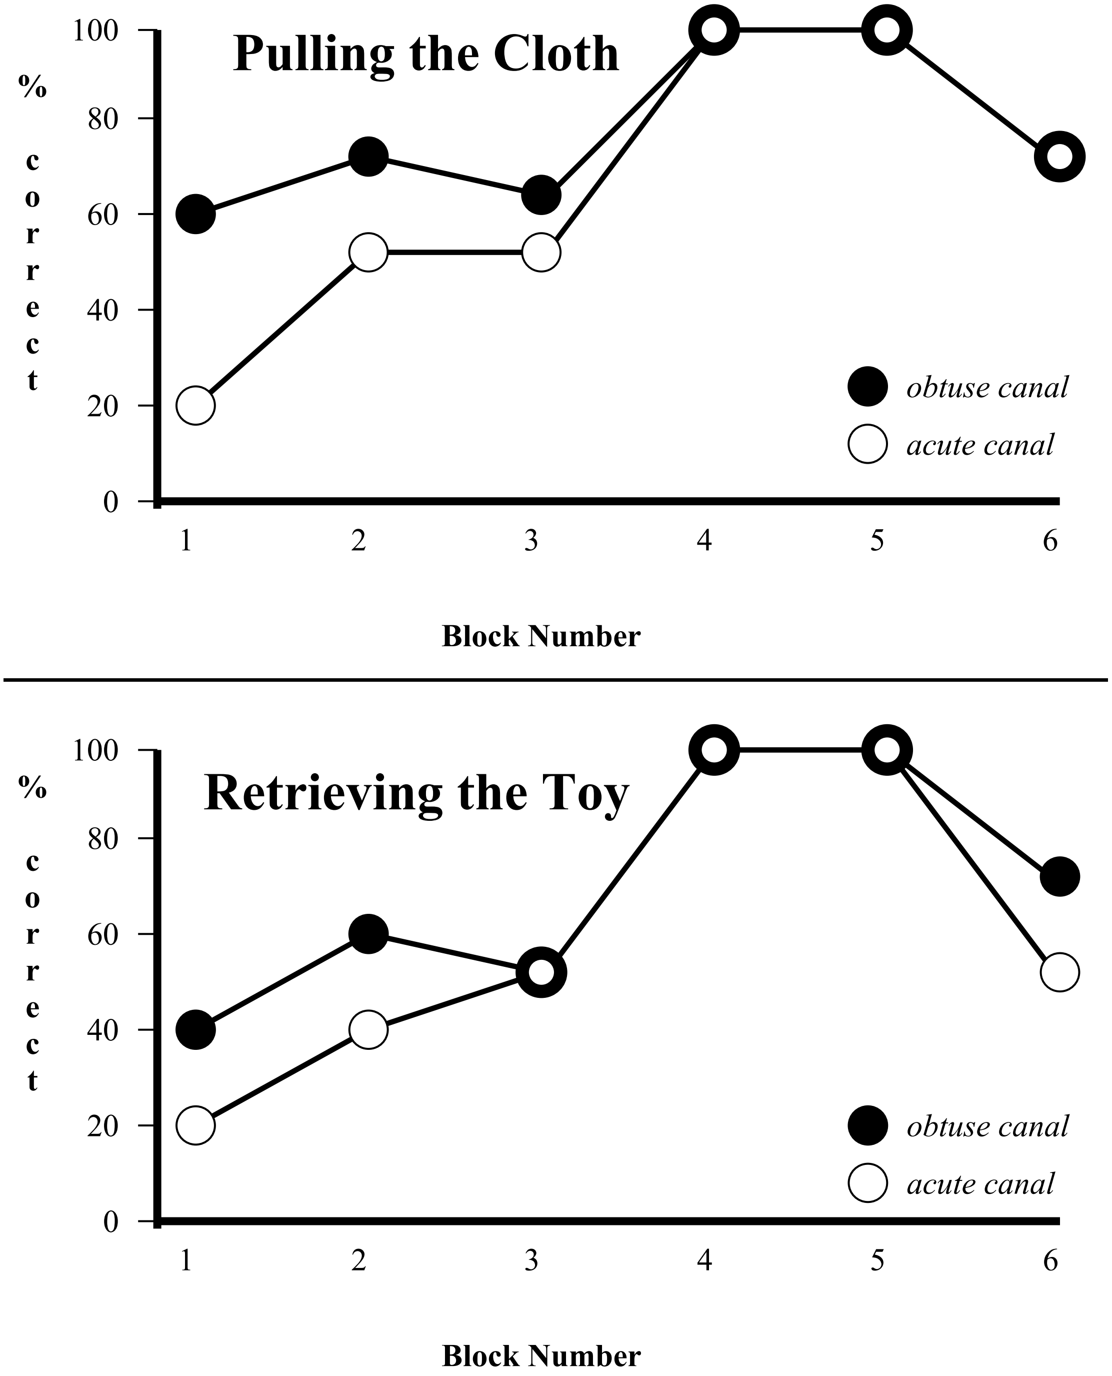

Figure C.1 shows the percentage of trials on which the average infant was successful during each of the six blocks. Though this is not a statistical test, we can see intuitively the suggestion that both hypotheses are plausible: domain general strategy development is gradual and infants can learn a domain general problem solving strategy in a particular context. The statistical tests that follow further indicate that the hypotheses are plausible. All statistical tests are one-tailed because a particular direction was always predicted in advance, and the opposite difference would be evidence that falsifies the hypothesis. See the result section for more complete explanations of the statistical tests.

Domain General Strategy Acquisition is Gradual

Figure C.1 suggests that, at first, the infant's performance on the obtuse canal was better than their performance on the acute canal. A paired sample t-test of acute versus obtuse trials in the first two blocks shows this difference is near marginally significant when performance is judged by pulling ($\underline{t}(4)=2.06$ $\underline{p}=.05$) or retrieving ($\underline{t}(4)=1.37$ $\underline{p}=.12$). These differences are consistent with the Gradual Strategy Acquisition Perspective (figure 3.2) but are inconsistent with the Complete Strategy Repertoire Perspective (figure 3.1).

Domain General Strategies Can be Learned

Figure C.1 suggests that infants' performance improves as they gain experience solving the canal problems. A correlation of success and block number for the first four blocks suggests that performance on acute canals improves during the experiment when performance is judged by pulling (r=.59 p<.01) or retrieving (r=.59 p<.01). These correlations are consistent with a Learning/Experience model, because experience solving problems predicts success, but these correlations are not consistent with an Innate/Constraint Model. The correlations of success and block number for the first four blocks was not significant for the obtuse canal when performance was judged by pulling (r=.23 p<.20) or retrieving (r=.33 p<.10). This may be because the 12- to 13-month-old infants began the experiment with successful performance on the obtuse canal trials.

Discussion

From this pilot study, the method was refined to limit discomfort to the infant. This study also suggests that it is reasonable to expect infants to be capable of completing the task. The results of this study provide preliminary evidence suggesting that both the milestone and the mechanism hypotheses are plausible.

References

Adolph, K. E., Eppler, M. A., & Gibson, E. J. (1993). Development of perception of affordance. Advances in Infancy Research. 8, 51-98.

Alibali, M. W. (1999). How children change their minds: Strategy change can be gradual or abrupt. Developmental Psychology, 35(1), 127-145.

Anderson, J. R. (1993). Problem solving and learning. American Psychologist, 48(1), 35-44.

Anderson, J. R. (1995). Cognitive psychology and its implications. (4th ed). USA: Freeman.

Anderson, J. R., Fincham, J. M., & Douglas, S. (1997). The role of examples and rules in the acquisition of a cognitive skill. Journal of Experimental Psychology, 23(4), 932-945.

Bahrick, L. E., Hernandez-Reif, M., & Pickens, J. N. (1997). The effects of retrieval cues on visual preferences and memory in infancy: Evidence for a four-phase attention function. Journal of Experimental Child Psychology, 67, 1-20.

Bates, E., Carlson-Luden, V., & Bretherton, I. (1980). Perceptual aspects of tool using in infancy. Infant Behavior and Development, 3, 127-140.

Bauer, P. J. & Hertsgaard, L. A. (1993) Increasing steps in recall of events: Factors facilitating immediate and long-term memory in 13.5- and 16.5-month-old children. Child Development, 64, 1204-1223.

Brown, A. L. (1989). Analogical learning and transfer: What develops? In S. Vosniagdou and A. Ortony (eds), similarity and analogical reasoning, 369-412, Cambridge: Cambridge University Press.

Brown, A. L. (1990). Domain-specific principles affect learning and transfer in children. Cognitive Science, 14, 107-133.

Brown, A. L. & Campione, J. C. (1984). Three faces of transfer: Implications for early competence, individual differences, and instruction. In M. Lamb, A. Brown, & B. Rogoff (Eds.). Advances in Developmental Psychology, 3, (pp. 143-192). Hillsdale: Erlbaum.

Brown, A. L. & Kane, M. J. (1988). Preschool children can learn to transfer: learning to learn and learning from example. Cognitive Psychology, 20, 493-523.

Brown, A. L., Kane, M. J., & Long, C. (1989). Analogical transfer in young children: Analogies as tools for communication and exposition. Applied Cognitive Psychology, 3, 275-293.

Bruner, J. S. (1973). Organization of early skilled action. Child Development, 44, 1-11.

Case, R. (1985). Intellectual development : Birth to adulthood, Orlando : Academic.

Case, R. (1991). The mind's staircase: Exploring the conceptual underpinnings of children's thought and knowledge. Hillsdale: Lawrence Erlbaum.

Chase, W. G. & Simon, H. A. (1973). The mind's eye in chess. In W. G. Chase (Ed.). Visual Information Processing. New York: Academic.

Chen, Z., Sanchez, R. P., & Campbell, T. (1997). From beyond to within their grasp: The rudiments of analogical problem solving in 10- and 13-months-olds. Developmental Psychology, 33(5), 790-801.

Chen, Z. & Siegler, R. S. (2000) Across the great divide: Bridging the gap between understanding of toddlers' and older children's thinking. Monographs of the Society for Research in Child Development, 65(2, Serial No. 261).

Cole, P. M., Michel, M. K., & Teti, L. O. (1994). The development of emotion regulation and dysregulation: A clinical perspective. Monographs of the Society for Research in Child Development, 59(2-3), 73-100, 250-283.

Diamond, A. (1985). Development of ability to use recall to guide action, as indicated by infants' performance on A-not-B. Child Development, 56, 868-883.

Dunker, K. (1945). On problem solving (translated by L. S. Lees). Psychological Monographs, 58, (Serial No. 270)

Eppler, Marion A., Adolph, K. E., & Weiner, T. (1996). The developmental relationship between infants' exploration and action on slanted surfaces. Infant Behavior and Development, 19, 259-264.

Frye, D. (1980). Stages of development: The stage IV error. Infant Behavior and Development, 3, 115-126.

Gardener, W. & Rogoff, B. (1990). Children's deliberateness of planning according to task circumstance. Developmental Psychology, 26 480-487.

Gilmore, R. O. & Johnson, M. H. (1995). Working memory in infancy: Six month olds' performance on two versions of the oculomotor delayed response task. Journal of Experimental

Child Psychology, 59, 397-418

Goswami (1996) Analogical reasoning and cognitive development. In H. Reese (Ed.), Advances in child development and behavior (Vol. 26). New York: Academic Press.

Grobman, K. H. (1998). Equal sign conceptual understanding and early algebra problem solving. Unpublished master's thesis, Carnegie Mellon University.

Haith, M. M. (1994). Who put the cog in infant cognition? Is rich interpretation too costly? Infant Behavior and Development, 21(2), 167-179

Henderson, B. B. and Dias, L. (1987). An exploratory study of infant problem solving in natural environments. Ethology and Sociobiology, 8, 205-213.

Karmiloff-Smith, A. (1986). Stage/structure versus phase/process in modeling linguistic and cognitive development. In I. Levin (Ed.)., Stage and structure: Reopening the debate. Norwood: Ablex.

Klahr, D. (1992). Information-processing approaches to cognitive development. In M. H. Bornstein, H. Marc., & M. E. Lamb (Eds.), Developmental psychology: An advanced textbook (3rd ed.). (pp. 273-335). Hillsdale: Lawrence Erlbaum.

Klahr, D. (1998). Discovering the present by predicting the future. In *EDITORS* (Eds.), Planning in Problem Solving, (pp. 176-217). *TOWN:PUBLISHER.*

Kotovsky, K. & Simon, H. A. (1990). What makes some problems really hard: explorations in the problem space of difficulty. Cognitive Psychology, 22, 143-183.

Mosier, C. E. & Rogoff, B. (1994). Infants' instrumental use of their mothers to achieve their goals. Child Development, 65, 70-79.

Michon, J. A. & Akyuerek, A. (1992). Soar: a cognitive architecture in perspective: A tribute to Allen Newell. Norwell, MA: Kluwer.

Munakata, Y., McClelland, J. L., Johnson, M. H., & Siegler, R. S. (1997). Rethinking infant knowledge: Toward an adaptive process account of successes and failures in object permanence tasks. Psychological Review. 104(4), 686-713.

Newel, A. (1980). Reasoning, problem-solving, and decision processes: The problem spaces as a fundamental category. In R. Nickerson (Ed.). Attention and performance (8). Hillsdale: Erlbaum.

Newel, A., & Simon, H. A. (1972). Human Problem Solving. Endlewood Cliffs: Prentice-Hill.

Piaget, J. (1954). The construction of reality in the child. New York: Basic Books.

Rabinowitz, M. F., Howe, M. L, & Saudners, K. (2002) Age, memory load, and individual differences in working memory as determinants of class-inclusion reasoning. Journal of Experimental Child Psychology. 81(2), 157-193.

Rittle-Johnson, B. & Alibali, M. W. (1999) Conceptual and procedural knowledge of mathematics: Does one lead to the other? Journal of Educational Psychology. 91(1), 175-189.

Ross, B. H., & Kilbane, M. C. (1997). Effects of Principle explanation and superficial similarity on analogical mapping in problem solving. Journal of Experimental Psychology. 23(2), 427-440.

Schauble, L. (1990). Belief revision in children: The role of prior knowledge and strategies for generating evidence. Journal of Experimental Child Psychology, 49, 31-57.

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

Siegler, R. S. (1983). Five generalizations about cognitive development. American Psychologist, 38(3), 263-277.

Siegler, R. S. (1998). Children's Thinking. (3rd ed.). Upper Saddle River: Prentice-Hill.

Siegler, R. S., & Shipley, C. (1995). Variation, selection, and cognitive change. in T. Simon and G. Halford (Eds.). Developing cognitive competence: New approaches to process modeling. (pp. 31-76). Hillsdale: Erlbaum.

Swartz, B. B., & Reznick, J. S. (1999). Measuring infant spatial working memory using a modified delayed-response procedure. Memory, 7(1), 1-17.

Sweller, J. (1988). Cognitive Load during problem solving: Effects on learning. Cognitive Science, 12(2), 257-285

Wellman, H. M. & Gelman, S. A., (1998). Acquisition of knowledge. In W. Damon (Series Ed.) & D. Kuhn & R. S. Siegler (Vol. Eds.), Handbook of child psychology, Vol.2: Cognition, perception, & language. (5th ed.). New York: Wiley.

Willatts, P. (1984). Stages in the development of intentional search by young infants. Developmental Psychology, 20(3), 389-396.

Willatts, P. (1984). The stage-IV infant's solution of problems requiring the use of supports. Infant Behavior and Development, 7, 125-134.

Willatts, P. (1985). Adjustments of means-ends coordination and the representation of spatial relations in the production of search errors by infants. British Journal of Developmental Psychology, 3, 259-272.

Willatts, P. (1999). Development of means-end behavior in young infants: Pulling a support to retrieve a distant object. Developmental Psychology, 35(3), 651-667.