

## COORDINATING MULTIPLE REPRESENTATIONS FOR REASONING ABOUT MECHANICAL DEVICES

**Daniel L. Schwartz**

Vanderbilt University

**Mary Hegarty**

University of California, Santa Barbara

### Abstract

People often draw inferences about the same mechanical situation using distinct reasoning methods. When reasoning about chains of gears, for example, people may simulate the motions of successive gears, or they may use a simple odd/even rule. We consider psychological and computational evidence as to when and why people use different reasoning methods to draw inferences about mechanical situations. We describe two types of reasoning strategies and their computational implementations. One strategy uses discrete rules to describe the behavior of a system, and the other uses analog simulations to depict the behavior of a system. We develop the strengths and weaknesses of each form of reasoning as a backdrop for considering the strategic knowledge that may guide people between the two forms of reasoning. We present a rough outline of an idealized strategic hierarchy for managing reasoning methods and memory resources. People sometimes follow this idealized hierarchy, but we present new evidence revealing situations in which people do not adaptively rely on an appropriate reasoning strategy. We conclude with a discussion of why it is useful to consider the meta-knowledge that brokers the use of different forms of representation.

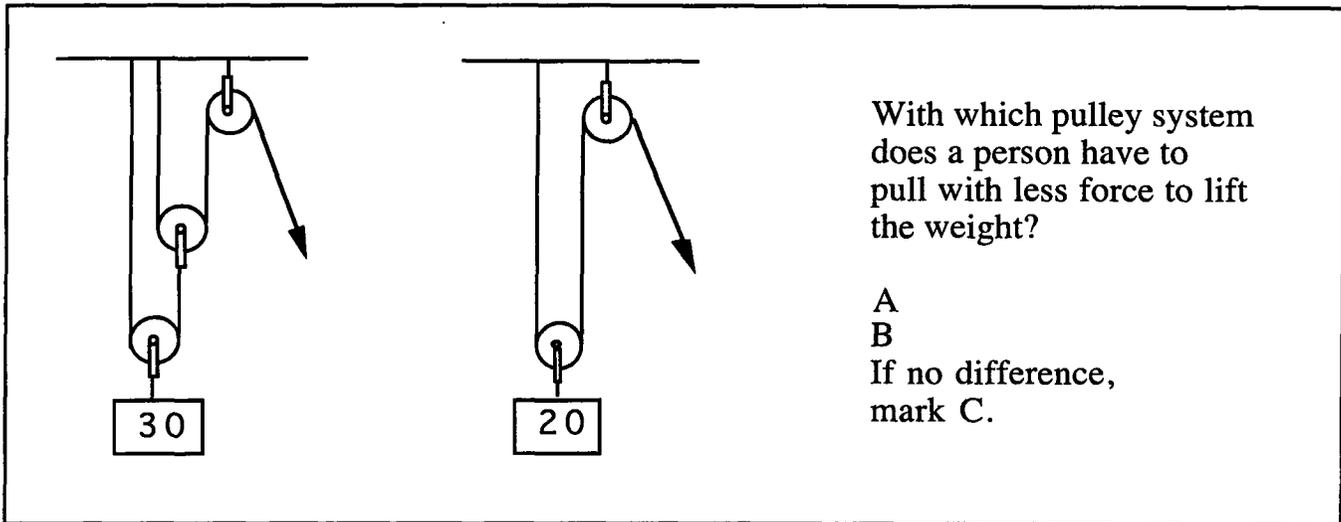
### Introduction

People arrive at conclusions about the same mechanical situation in many different ways. For example, imagine trying to determine the motion of the last gear in a five gear chain when the first gear turns clockwise. On one occasion, a person might imagine the behavior of each successive gear in the chain. On another occasion, the same person might apply a simple odd/even rule which states that odd-numbered gears will turn in one direction and even-numbered gears will turn in the opposite direction (Schwartz & Black, in press a). As a result, it has been

proposed that making inferences about the mechanical systems involves hybrid reasoning processes (Narayanan, Suwa & Motoda, 1995). Our goal in this presentation is to consider psychological and computational factors that determine how an individual chooses to solve a problem. We will begin by describing two classes of computer model that capture distinct ways people reason about a mechanical situation. One class of model uses rule-based reasoning methods to infer a description of the behavior of a device (Hegarty, Just & Morrison, 1988; Hegarty, 1992), whereas the other uses an analog simulation to depict the behavior of a device (Schwartz & Black, in press b). We will develop the strengths and weaknesses of each form of reasoning. This analysis sets the stage for considering strategic knowledge that may guide people between the two forms of reasoning. We present a rough outline of an idealized strategic hierarchy. We then consider empirical evidence as to whether people follow this idealized hierarchy. We conclude with a discussion of why it is useful to consider both multiple forms of representation and the meta-knowledge that brokers their use.

### Rule Based Models

One type of computer simulation describes how rule-based reasoning can be applied wholesale to classes of well-known problems. Or, in the case where the total system behavior is unknown, the rules can be used to decompose the structure of a device into a sequence of component interactions that are known. Empirical evidence for this type of rule-based reasoning comes from studies in which people viewed depictions of two different pulley systems, as shown in Figure 1, and made judgments about which system would require less force to lift a given weight (Hegarty, Just & Morrison, 1988). Verbal protocols revealed that subjects made their judgments based on rules of mechanical reasoning, such as "a system with more pulleys requires less force". The data also showed systematic individual differences among the subjects.



**Figure 1.** Sample pulley problem (Hegarty, Just & Morrison, 1988)

High-ability subjects were more able to identify which attributes of a mechanical system are relevant to its mechanical function, used rules more consistently, and were better able to quantitatively combine information about two or more relevant attributes.

Subjects' reasoning was modeled using a production system implemented in SOAR. The productions decomposed the mechanical devices into elementary components (pulleys, ropes etc.) and compared the devices on the basis of the number of each type of component that the devices contained. For example, in solving the problem in Figure 1, the simulation might note that pulley system A has 3 pulleys and pulley system B has 2 pulleys. On the basis of the rule "a system with more pulleys requires less force" the model would then hypothesize that pulley system A requires less force. Individual differences in mechanical reasoning were modeled by giving the production system different mechanical rules and different preferences among the rules that determined which answer was chosen for a problem when mechanical rules produced conflicting hypotheses.

In another set of studies, people verified how a given component of a pulley system would move when the free end of the rope was pulled, as in the problem presented in Figure 2 (Hegarty, 1992). Because this task requires reasoning about device kinematics, Hegarty referred to one piece of the reasoning as "mental animation." Hegarty modeled the reasoning for this kinematic task using a production system. In this model, the simulation also decomposed the system into elementary components, but also kept track of the spatial relations between the components. Static information about the spatial relations between two components, and kinematic information about the movement of one of the two components, served as the

conditions for each production rule. The action of a production rule was to infer the motion of the other component. For example, the following production from Hegarty's model would infer that the rope over the upper pulley in Figure 2 is moving to the right.

IF a rope strand (RS1) lies over a pulley,  
and it is attached to another rope strand (RS2) on  
the right,  
and RS2 is moving down  
THEN infer that RS1 is moving to the right over the  
pulley

Note that the mechanical rules in this model could not be applied wholesale to a problem as in the Hegarty, Just & Morrison model. Rather, the simulation inferred the motion of the components, one by one, following the causal chain of events. This model accounted for reaction times and eye fixations as people inferred the motion of different components in the pulley devices. For example, both the model and the human subjects took longer to infer the motion of components later in the causal chain. The causal ordering of inferences was also consistent with data showing that people have great difficulty making inferences against the causal chain of events in a system (Hegarty, 1992).

Further research has indicated that differences in mental animation performance can be predicted by spatial visualization ability (Hegarty & Sims, 1994) and that mental animation interferes more with a concurrent spatial memory load than a concurrent verbal memory load (Sims & Hegarty, in press). The foregoing results suggest that kinematic inferences about the interaction of adjacent device components rely on a working memory that is specialized for processing spatial information. Work in

When the free end of the rope is pulled, the middle pulley turns counterclockwise

- True
- False

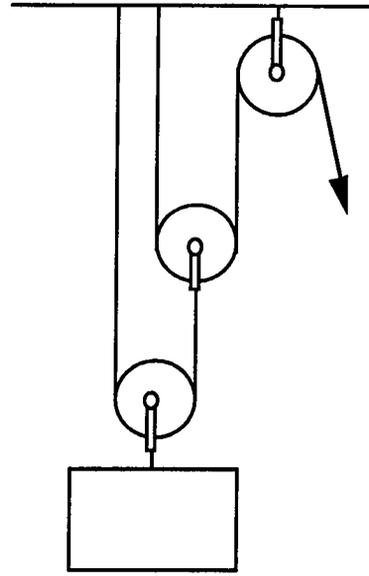


Figure 2. Sample mental animation problem (Hegarty 1992).

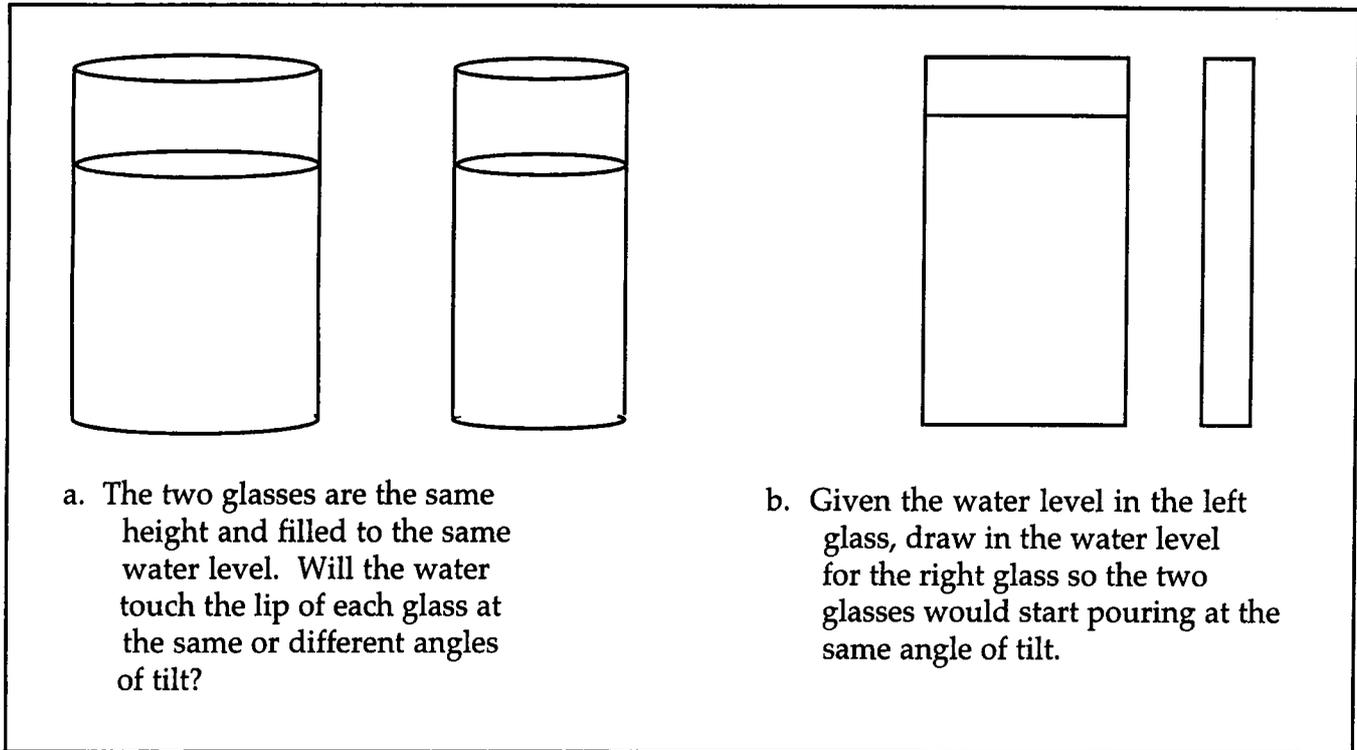
progress is focused on modeling individual differences in the mental animation task in terms of memory resources and strategies for managing those resources. For this work, Hegarty's mental animation model is being implemented in the 3CAPS architecture (Just, Carpenter & Hemphill, 1995). 3CAPS is a production system architecture designed to model working memory constraints in complex cognitive tasks such as language comprehension and spatial inference. This architecture allows the user to specify pools of working memory resources for processing different types of information (e.g., a verbal and a spatial working memory pool) and to limit the amount of resources available to each pool. Differences in mental animation performance by high- and low-spatial subjects might be modeled by allocating more spatial working memory resources to the simulation of the high-spatial subjects, or by giving this simulation better heuristics for managing limited resources.

Although 3-CAPS works well for describing working memory constraints and uses, it follows a rule-based architecture that may underestimate the uniqueness of spatial reasoning. In informal verbal protocol studies of the mental animation task, subjects never articulated rules such as the sample production rule stated above. In contrast, their inference of the motion of pulley-system components was often accompanied by gestures imitating these component motions. Inferences that occur in spatial working memory may manifest themselves as analog spatial imagery, rather than as a set of discrete condition-action pairs.

### Depictive Models

The construct of a depictive model was developed to explain how people arrive at physical inferences through dynamic, analog imagery, and why some inferences can only be drawn through analog imagery (i.e., mental depiction, as we call it for physical inferences). Schwartz and T. Black (under review) demonstrated that most people can draw certain physical inferences only through mental depiction. In a series of experiments, subjects were shown the two glasses schematized in Figure 3a. A line drawn on each glass indicated the level of some imaginary water in the glass and the levels were the same for the two glasses. In some conditions (the static conditions) subjects were asked whether the imaginary water in the two glasses would reach the rims at the same or different angles of tilt. In these conditions, subjects rarely described the correct answer. In the other (dynamic) conditions, subjects closed their eyes and tilted (or imagined tilting) each glass until they "saw" that the pretend water had reached the rim. In these conditions, almost every subject correctly tilted a narrower glass further than a wider glass. The dynamic imagery of tilting the glasses supported the correct inference whereas non-dynamic, explicit solution attempts did not.

Interestingly, even after correctly tilting the glasses, subjects still made incorrect inferences when asked explicitly about the relative tilts of the two glasses. In fact, when shown the two glasses in Figure 3b and asked to



**Figure 3:** Versions of the water-pouring task (Schwartz & T. Black, under review).

indicate a water level for the thin glass that would lead to equivalent tilts for the two glasses, subjects consistently expressed the exact opposite of the true relationship (i.e., they thought the water would have to be further from the rim in the thin glass than the wide glass). This lack of a correct qualitative rule that could operate independently of analog imagery suggests one psychological constraint on a computer model of mental depiction. That is, such a computer model should not rely on a rule that first describes system behaviors and then subsequently manipulates spatial representations of the system to fit this description (e.g., Kosslyn, 1980). Instead, the behavior of the system must arise as a result of the analog transformation itself.

Further research demonstrated that the mental depiction supported a dynamic rather than kinematic solution to the problem. Subjects did not solve the problem solely on the basis of shape, but rather also depended on a model of forces in action. So, for example, when the problem changed physical constraints but not shape relationships (e.g., by presenting the glasses upside down or by placing weights on the bottom of the glasses), this influenced the subjects' imagery. Thus, a second property of a computer simulation of mental depiction is that it should represent physical knowledge in an analog form that responds to dynamic as well as kinematic relations.

Another series of experiments developed further evidence regarding the characteristics of mental depiction. In these experiments (Schwartz & Black, in press b), subjects had to determine whether the knob and groove on two gears would meet if the gears were rotated inward. Figure 4 provides an example of the task. Subjects' latencies were in direct proportion to the distance they would have to imagine rotating the gears to bring the knob and groove into alignment, providing evidence of analog imagery (cf. Shepard & Cooper, 1986). Interestingly, subjects showed distinct latency patterns depending on the visual presentation of the gears. When the gears were presented as realistic 3-D drawings, solution times were dependent on the circumferential disparity (arc lengths) of the marks from the alignment point. Subjects imagined the gears rotating point-for-point along their circumferences, as though "mental friction" were keeping the tandem rotation rates coordinated. (We use the term "mental friction" as short hand for describing how surfaces and forces interact.) However, when the gears were presented as simple line drawings, solution times were dependent on the angular disparity of the marks from the alignment point. In this case, subjects appeared to have first determined the appropriate rate of rotation for each gear using a rule (e.g., twice as big means half as fast) and then set them into motion. Similar results were found when

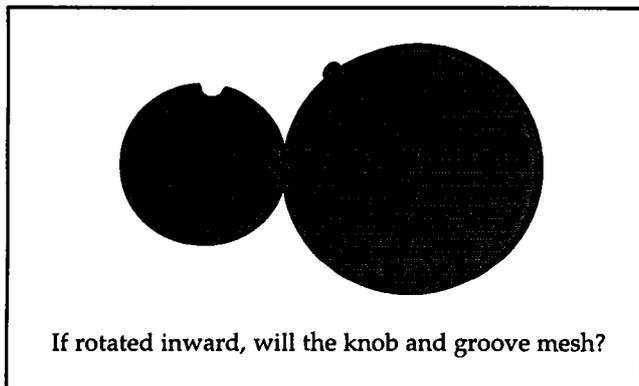


Figure 4. Two gear task (Schwartz & J. Black, in press b)

subjects were asked to reason about realistically or schematically drawn hinges (Schwartz, 1995).

One explanation for the effect of realism is that the realistic drawings provided important surface information that was not available in the simple line drawings. To coordinate the physical interaction of the gears, it may be necessary to model surfaces that can communicate forces. The line drawings might not have provided the subjects with sufficient surface information to mentally regulate the interactions of the gears. This suggests two further features for a computer simulation of mental depiction. First, the simulation should model information access both from memory and from the external stimuli. Second, the simulation should require sufficient information to yield determinate imagery. This contrasts with another class of models of mechanical inference, such as de Kleer & Brown's (1984) model, in which the program generates a space of possible solutions, all consistent with partial information, from which it then selects heuristically.

An object-oriented model of mental depiction for the two gear task was developed to meet the four requirements laid out above. The model had the following characteristics:

(1) The simulation did not rely on rules that pre-specified the trajectory of a transformation. Instead the course of the transformation was determined by constraints built into object-based representations of the gears. Data and procedure encapsulation within the gear objects ensured that the program could not extract attributes of the gear objects (e.g., a numerical size) to draw a rule-based inference about the appropriate relative turning rates.

(2) The objects embodied analog representations that modeled kinematic and dynamic relations. The gears were represented by a set of vectors. As in other spatial representations (Marr, 1982), these vectors indicated surface edges and curvature. The vectors were further augmented so that they also represented, in analog form, physical properties such as friction, rigidity, and direction of force. When one of the gear representations was

perturbed, each of its spatial and physical vectors sent a message that it had changed. If the second gear overlapped the extended local space of the first gear, some of these messages were caught by the surface vectors of the second gear. The constraints on the vectors of the second gear interacted with the motion vectors of the first gear to determine how it should respond. This occurred in parallel across the affected surface of the second gear.

(3) The simulation used the inheritance hierarchy of the object-oriented paradigm to model retrieval and object construction. Object construction was conceived of as the retrieval or perception of a collection of constraints that were constructed into an object. When the objects were fully instantiated, simple primitive constraints retrieved from the hierarchy combined and interacted in novel ways to reveal higher order patterns of behavior. The idea that people can rely on relatively low-level, imagined constraints to infer higher order system behaviors is consistent with psychological theories of spatial cognition. Shepard (1994), for example, has proposed a constraint-based geometric kinematics such that basic constraints on imagined rotations and translations yield, in combination, more complex transformations such as a screw or pivot displacement. The novel psychological claims here are twofold. First, imagery constraints may also capture physical regularities like gravity or the slowing effects of friction on two moving surfaces (Freyd, 1987; Hubbard, 1995). Second, the constraints are learned and are therefore subject to retrieval cues and mechanisms. For example, when the simulation received the line drawings rather than the realistic drawings, it could not retrieve constraints on the behavior of a physical surface because there was no surface information in the realistic drawings to cue this knowledge.

(4) The program required sufficient information to infer a determinate representation of object motion. Psychological research has shown that spatial working memory representations are determinate with respect to spatial relations (cf. Mani & Johnson-Laird, 1982). For example, people can relatively easily retrieve from long-term memory any perspective within a room, but once they have that determinate perspective in mind, it becomes very difficult to imagine other perspectives within the room (Rieser, Garing & Young, 1994). That is, representations that are indeterminate with respect to perspective in long-term memory become determinant in working memory. In the current case of a mechanical inference, the program required determinate physical information as well as spatial information. When the program received the realistic drawing it had sufficient surface information to determine the interactions of physical surfaces (i.e., point-for-point behavior). However, when the simulation received the line drawings, it could not infer the behaviors of the gears depictively because the properties of the gear surfaces were

indeterminate (e.g., they might have been non-rigid).

The simulation matched human behavior for the gear task. When it received the realistic drawing of the gears, the number of message passing cycles was proportional to the circumferential disparity of the marks. The set of constraints that modeled the point-for-point interaction of the gears also modeled the jamming of the knob into the side of the left hand gear when the marks would not align. Latency data from human subjects indicated that they also terminated their solutions when the knob would have jammed the two gears. The depictive simulation, however, was unable to model the line drawings. In this case, we augmented the model to include attribute extraction procedures (to find the numerical sizes of the gears) and global rules that could use the attributes to pre-determine and guide the relative angular velocities of the gears.

The computer simulation of mental depiction is similar to those of Funt (1980) and Gardin & Meltzer (1989). In both of these simulations, space is represented as an array structure that is intended to be an analog of perceived space. Physical knowledge takes the form of local constraints on object behaviors. The constraints regulate on-going object transformations within the spatial array. These constraints yield determinate transformations rather than a problem space of alternative transformations among which the program must choose. The differences between these models and ours involve issues of memory retrieval and the representation of space (i.e., the model of mental depiction does not rely on a global coordinate system in which the objects reside). However, for the current purposes, the important characteristic shared by these models and ours is that, unlike rule-based inferences which can derive solutions without changing a model of the system (e.g., odd gears turn the same direction), the knowledge embodied in constraints only has an effect during the course of an analog transformation. Thus, these models provide one class of explanation for how people could have powerful inferential knowledge that does not depend on prior rule and attribute abstraction from the objects of experience

### **Comparison of Rule-Based and Depictive Models**

Rule-based and depictive inferences have several complementary strengths and weaknesses. On the plus side, rule-based inferences are relatively quick, make few resource demands, have no limits on the complexity of systems that they can handle, apply to broad classes of well-defined situations, can reach provisional conclusions in low information situations, and permit higher-order reasoning (e.g., proportional reasoning) about device attributes. On the negative side, rules depend on either abstraction from experience or explicit instruction and

often have constraints on their applicability. Consequently, people might not have rules that they can readily apply to novel situations, or people might apply rules inappropriately in novel situations. Novices, in particular, might have abstracted incorrect attributes or applicability conditions from prior experience. For example, when asked to infer the motion of pulleys connected by ropes, some novices apply a rule that every other pulley turns the same direction. This can be interpreted as an overgeneralization of a rule that is applicable to chains of gears.

In contrast to rule-based models, depictive models are able to simulate novel devices; they can infer basic causal interactions by revealing the effects of novel combinations of interacting constraints. They can represent physical and spatial properties in the same analog representation, and therefore do not depend on prior abstractions that have converted perceptual experiences into discrete rules. Depictive models also have several limitations. They are deterministic such that they require sufficient specificity that they will reach a single outcome. They are resource intensive, and therefore, can only model a limited range of motion of two or three related objects at a time. They are constrained to operate in a manner analogous to perceptual experience, and thus, cannot operate over abstract descriptions of system attributes (e.g., numerical measures). While their results may guide behavior, depictive results may not always be readily available for redescription into a verbal form.

### **Strategy Choice in Mechanical Reasoning**

Given the complementary advantages and disadvantages of rule-based and depictive models, we propose a decision hierarchy for choosing an optimally efficient method for reasoning about a mechanical device. This hierarchy is not meant to cover all possible methods of reasoning, such as using an analogy. Rather, it is more circumscribed in its emphasis on directly applicable rules or depictions. In the rough, the reasoner should first look for a rule that can be applied wholesale to the device. If no such rule exists, the reasoner should attempt to decompose the mechanical device into a sequence of component interactions amenable to rule-based inferences. Finally, if no rules apply to a component interaction, the reasoner should use a mental depiction to simulate the behavior of coupled components. This portion of the hierarchy guides the access and use of relevant knowledge from a long term store.

To be complete, the hierarchy should also provide decision rules for managing working memory limitations. These depend on meta-knowledge of how much complexity can be simultaneously modeled, and strategic knowledge for how to store intermediate results when a system has been decomposed. With respect to the first

point, the imagery of a mental depiction cannot maintain sufficient resolution when the number of components becomes large (Kosslyn, 1980). For example, it is not possible to imagine the motion of all components of even a simple pulley system moving at once (Hegarty, 1992). Consequently, the system needs some knowledge about imagery limitations to be able to decompose a mechanism to an appropriate granularity. With respect to storing intermediate results, reasoners who make an inference about a sub-component of a device and do not immediately propagate the result to the next sub-component, must store the results for later inference. The form of storage interacts with the reasoning strategies at play (Baddeley, 1986; Logie, 1995). When a person uses a depiction to model a pair of interacting components, this presumably draws heavily on spatial working memory. Consequently, it would be inadvisable to store the results of a previous inference in spatial working memory when a person is about to make a new depictive inference. Instead, the reasoner might recode depictive results into a verbal code, or use an external memory store (i.e., draw an arrow on a diagram of a mechanical system to indicate the inferred direction of motion of a component). Thus, a second aspect of the decision hierarchy considers how to optimize the storage of intermediate results based on known and hypothesized capacity limitations.

To what extent do human inferences about mechanical systems follow this idealized hierarchy? In our research we have observed situations in which people are adaptive in choosing a strategy for mechanical reasoning. For example, when mentally animating a pulley system, people take account of the fact that depictive reasoning is resource intensive. Thus, they first decompose the device into a more manageable series of component interactions (Hegarty, 1992). As another example, when people first reason about a novel problem involving gear chains, they gesturally depict pairs of gears within the larger chain. After a few depictions, people abstract an odd/even rule from their depictions and quickly switch to the less resource-intensive, rule-based approach. However, if their rules subsequently fail in a novel situation, they fall back to a depictive strategy (Schwartz & Black, in press a). Finally, we have observed that when some people make a chain of inferences about subcomponents of a complex device, they recode their intermediate depictive results into a verbal format or an external representation, thus freeing up working memory resources to make further depictive inferences.

We have also observed situations in which people do not follow our idealized hierarchy in their mechanical reasoning strategies. For example, in one experiment, people were given the opportunity to make notes on diagrams of mechanical systems as they inferred the motion of system components. Some subjects failed to

take advantage of the opportunity to make notes, although note-making enhanced the performance of those who did take advantage of this opportunity (Hegarty & Steinhoff, 1994). In another experiment (Schwartz, in preparation), we asked several people to reason about the direction a fan would turn if a wind blew through it from behind. Some subjects reported decomposing the problem so they could reason about the behavior of one fan blade and solved the problem successfully. Others reported trying to imagine the behavior of the fan wholesale. Most of these failed at the task, presumably because this wholesale imagining overloads the resources of spatial working memory.

A recently completed experiment on a version of the water-pouring task (depicted in Figure 3a) provides another example of a failure of subjects' meta-knowledge of mechanical reasoning strategies. In one condition of the experiment, we observed whether subjects spontaneously adopt a depictive strategy or a rule-based strategy in solving the water-pouring task. Fifteen subjects were shown empty glasses similar to those in Figure 3a and had to predict whether the glasses should be tilted the same or different amounts for the pretend water to just reach the rim. Twelve out of the fifteen subjects reached for and manipulated the glasses and imaginary water prior to responding verbally. This spontaneous adoption of a depiction strategy is consistent with previous research showing that people rely on mental depiction in situations of novelty (Schwartz & Black, in press a).

In the second condition of this experiment we assessed whether people trust their mental depictions or their verbal judgments in a situation where the two are in conflict. Fifteen subjects were asked to close their eyes and tilt each glass in turn until they imagined that the pretend water had just reached the rim of the glass. They tilted each glass four times using a different imagined water level each time. As expected, for each of the four water levels, the subjects tilted the thin glass further than the wide glass [ $F(1,14)=68.5, MS_e = 29.5, p<.001$ ], and there was a positive monotonic relationship between the distance of the water level from the rim and how far they tilted the glasses [ $F(3,42)=104.7, MS_e = 23.6, p<.001$ ]. After tilting all eight glasses, we asked the subjects whether the thin and wide glasses would need to be tilted the same or different amounts. As in previous research (Schwartz & T. Black, under review), subjects' verbal judgments were inaccurate - only two of the fifteen subjects correctly stated that a thin glass would need to be tilted further than a wide glass.

The critical data came from the next question we asked the thirteen subjects who made the incorrect verbal judgment. We told the subjects of the inconsistency between their tilts and their verbal judgments. That is, we told them that they had consistently tilted the thin glass further than the corresponding wide glass, although this was inconsistent with their verbal judgments. We then

asked them whether they believed their verbal response or the evidence from their imagery. Eleven of the thirteen subjects incorrectly chose their verbal response. Thus, although 80% the subjects in the first condition spontaneously manipulated the empty glasses to help reason about the task, 84.6% of the subjects in the second condition reported that they would not trust the results of this strategy [ $\chi^2(1, N=28)=11.6, p<.001$ ]. The picture that emerges from these data is that people adaptively rely on their imagination when confronting a novel problem (i.e., they reach and manipulate the glasses), but will disregard this evidence if they have previously made a rule-based judgment that contradicts the imagery evidence. This result demonstrates the complexity of people's beliefs about their own strategies.

### Where Artificial Intelligence and Psychology can Meet

We suggest three potential lines of research for profitably considering how the two distinct reasoning systems interact. First, by examining when people move between different forms of reasoning, it may be possible to eliminate degrees of freedom in the description of these reasoning types. That is, a computer implementation with explanatory power should be able to account for the "when" of a representation, as well as, the "what." Switches between the two forms of reasoning may inform us about the conditions of applicability of each form.

Second, it is worthwhile to investigate and describe meta-knowledge about memory resources and the utility of different representations. People, for example, must have some knowledge for deciding which way to represent a problem initially and for switching to a new representation because a particular reasoning approach is failing. Similarly, people must have knowledge about their own memory limitations. For example, they often draw a directional arrow on a mechanical component, presumably because they know that the inferred motion of the component will be wiped out upon the next inference. Although people often times show adaptive strategy decisions, they also demonstrate non-adaptive choices. Therefore it is important to combine studies of people's meta-knowledge with empirical studies of the actual strengths and limitations of the different reasoning methods. We suspect that failure on many mechanical reasoning tasks is not due to a lack of relevant knowledge about the physical interaction. Rather, it may be due to a failure to adaptively apply the most effective solution method because of limited meta-knowledge.

Third, it seems worthwhile to take a prescriptive approach by asking what pieces of meta-knowledge, what sensitivity to one's own reasoning success, and what awareness of memory limitations would improve an

individual's mechanical reasoning. Even if people do have a set of defaults for moving through their different reasoning approaches, we presume these defaults are learned, and are teachable. This is an ideal application of AI to cognitive psychology. Decision hierarchies such as the one described above could be implemented as computer models. One could simulate different possible decision trees to see which ones allow the two model-based reasoning methods to most optimally complement one another. The results of these simulations might provide excellent guidance for the design of experiments and instruction that could provide for an empirical test of how people coordinate multiple representations and reasoning methods.

### References

- Baddeley, A. D. 1986. *Working Memory*. New York: Oxford University Press.
- de Kleer, J. & Brown, J. S. 1984. A qualitative physics based on confluences. *Artificial Intelligence*, 24, 7-83.
- Freyd, J. J. 1987. Dynamic mental representations. *Psychological Review*, 94, 427-438.
- Funt, B. V. 1980. Problem-solving with diagrammatic representations. *Artificial Intelligence*, 13, 201-230.
- Gardin, F. & Meltzer, B. 1989. Analogical representations of naive physics. *Artificial Intelligence*, 38, 139-159.
- Hegarty, M. 1992. Mental animation: Inferring motion from static diagrams of mechanical systems. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 18, 1084-1102.
- Hegarty, Just & Morrison 1988. Mental models of mechanical systems: Individual differences in qualitative and quantitative reasoning. *Cognitive Psychology*, 20, 191-236.
- Hegarty, M. & Sims, V. K. 1994. Individual differences in mental animation during mechanical reasoning. *Memory & Cognition*, 22, 455-473.
- Hegarty, M. & Steinhoff, K. 1994. Use of diagrams as external memory in a mechanical reasoning task. Paper presented at the *Annual meeting of the American Educational Research Association*, New Orleans, April.
- Hubbard, T. L. 1995. Cognitive representation of motion: Evidence for friction and gravity analogues. *Journal of Experimental Psychology: Learning, Memory, and*

*Cognition*, 21, 241-254.

Just, M. A., Carpenter, P. A. & Hemphill, D. D. 1995. Constraints on processing capacity: Architectural or implementational? In D. Steier & T. Mitchell (Eds.), *Mind matters: A tribute to Allen Newell*. Hillsdale, NJ: Erlbaum.

Kosslyn, S. M. 1980. *Imagery and mind*. Cambridge, MA: Harvard University Press.

Logie, R. H. 1995. *Visuo-spatial working memory*. Hove, UK: Erlbaum.

Mani, K. & Johnson-Laird, P. N. 1982. The mental representation of spatial descriptions. *Memory & Cognition*, 10, 181-187.

Marr, D. 1982. *Vision*. San Francisco, CA: W. H. Freeman & Co.

Narayanan, N. H. , Suwa, M. & Motoda, H. 1995. Behavior hypothesis from schematic diagrams. In J. Glasgow, N. H. Narayanan & B Chandrasekaran (Eds.), *Diagrammatic Reasoning: Cognitive and Computational Perspectives*. Menlo Park, CA.: AAAI Press.

Rieser, J. J., Garing, A. E. & Young, M. F. 1994. Imagery, action, and young children's spatial orientation: It's not being there that counts, it's what one has in mind. *Child Development*, 65, 1262-1278.

Schwartz, D. L. in preparation. *Some strategic determinants of mechanical reasoning ability*. Manuscript in preparation.

Schwartz, D. L. 1995. Reasoning about the referent of a picture versus reasoning about the picture as the referent: An effect of visual realism. *Memory & Cognition*, 23, 709-722.

Schwartz, D. & Black, J. in press a. Shuttling between depictive models and abstract rules: Induction and fallback. *Cognitive Science*.

Schwartz, D. & Black, J. in press b. Analog imagery in mental model reasoning: Depictive models. *Cognitive Psychology*.

Schwartz, D. & Black, T. under review. *Situated physical inference: Depictive knowing versus descriptive beliefs*. Manuscript submitted for publication.

Shepard, R. N. 1994. Perceptual-cognitive universals as reflections of the world. *Psychonomic Bulletin and*

*Review*, 1, 2-28.

Shepard, R. N. & Cooper, L. A. 1986. *Mental images and their transformations*. Cambridge, MA: The MIT Press.

Sims, V. K. & Hegarty, M. in press. Mental animation in the visual-spatial sketchpad: Evidence from dual-task studies. *Memory & Cognition*.