

Mental Imagery in Problem Solving: An Eye Tracking Study

Daesub Yoon and N. Hari Narayanan
Intelligent and Interactive Systems Laboratory
Department of Computer Science and Software Engineering
Auburn University, Auburn, AL, 36849, U.S.A
{yoondae, narayan}@eng.auburn.edu

Abstract

Cognitive models and empirical studies of problem solving in visuo-spatial and causal domains suggest that problem solving tasks in such domains invoke cognitive processes involving mental animation and imagery. If these internal processes are externally manifested in the form of eye movements, such tasks present situations in which the trajectory of a user's visual attention can provide clues regarding his or her information needs to an Attentive User Interface [Vertegaal 2002]. In this paper, we briefly review research related to problem solving that involves mental imagery, and describe an experiment that looked for evidence and effects of an imagery strategy in problem solving. We eye-tracked 90 subjects solving two causal reasoning problems, one in which a diagram of the problem appeared on the stimulus display, and a second related problem that was posed on a blank display. Results indicated that 42% of the subjects employed mental imagery and visually scanned the display in a correspondingly systematic fashion. This suggests that information displays that respond to a user's visual attention trajectory, a kind of Attentive User Interface, are more likely to benefit this class of users.

CR Categories: H.1.2 [Models and Principles]: User/Machine Systems -- Human information processing; H.5.2 [Information Interfaces and Presentation]: User Interfaces -- Theory and methods.

Keywords: empirical study, mental imagery, reactive information display, attentive user interface.

1 Introduction

Displays that track the user's attention and react to present the right information at the right place and in the right time is an emerging frontier in Human-Computer Interaction research. While work on this topic has at least a fourteen-year history [e.g. Starker and Bolt 1990], advances have been sporadic. One reason may be a lack of models and principles for predicting a user's attention trajectory and information needs. Cognitive models [Narayanan and Hegarty 2002], computational models [Narayanan and Chandrasekaran 1991], and empirical studies

[Hegarty 1992] of problem solving in visuo-spatial and causal domains suggest that problem solving tasks in such domains invoke cognitive processes involving mental animation and imagery. If these internal processes are externally manifested in the form of eye movements, such tasks present situations in which the trajectory of a user's visual attention is systematic and predictable, and which, therefore, can provide clues regarding his or her information needs. If so, it opens up an interesting class of application areas for Reactive Information Displays [Narayanan & Yoon, 2003], which form a subclass of Attentive User Interfaces [CACM 2003].

In this paper we briefly review research related to problem solving involving mental animation and imagery, discuss its relevance to the design of information displays that track and react to the user's visual attention shifts, describe an experiment that looked for evidence and effects of mental imagery in problem solving, and describe its implications for future research on Reactive Information Displays. In the experiment we eye-tracked 90 subjects solving two causal reasoning problems, one in which a diagram of the problem appeared on the stimulus display, and a second related problem that was posed on a blank display. The purpose was to categorize users into those using, and not using, an imagery strategy, and to compare their performance in terms of several dependent measures. Results indicated that 42% of the users engaged in imagery during problem solving. While this reasoning strategy did not influence accuracy, these users generally spent more time on the task, and visually scanned the display in a more systematic fashion. This suggests that there is a class of users who are likely to benefit from displays that track and react to their visual attention in order to dynamically provide additional information.

The rest of the paper is structured as follows. Section 2 discusses a class of domains and problems in which mental animation and imagery play a role in reasoning. The following section explains the notion of Reactive Information Displays. Section 4 provides a brief literature review of mental animation, imagery and a cognitive model that formed the basis of our research. Sections 5-8 present details of the experiment and its results. Section 9 concludes the paper with a discussion of the implications of our findings.

2 Visuo-Spatial and Causal Domains

Interactive multimedia information displays are increasingly being used to aid planners, decision-makers and problem solvers in domains as varied as weather forecasting, military strategy planning, and controlling complex systems like power plants. In this kind of task, the user is presented with information on domain objects that are both spatially distributed and causally related. The situation represented by these objects is typically dynamic as well, i.e. it evolves over time. The user has to comprehend the presented information and subsequently carry out a problem solving task such as explaining, planning, predicting or troubleshooting.

These kinds of domains share certain characteristics:

- Constituent systems of these domains exhibit hierarchical structures composed of subsystems and components.
- Subsystems and components exhibit natural behaviours or engineered functions.
- These component/subsystem behaviours causally influence (trigger, modify or prevent) behaviours of other components/subsystems.
- The propagation of these causal influences creates chains of events in the operation of the overall system, and gives rise to its overall behaviour and function.
- These chains of events extend in temporal and spatial dimensions.

Understanding and being able to predict these event chains is a critical aspect of problem solving in such domains. One example of this kind of problem solving is reasoning about the operation of mechanical devices from cross-sectional diagrams. For example, Figure 1 shows the cross section of a common mechanical device, a flushing cistern. This system consists of two subsystems: a water outlet system and a water inlet system. The components of the former are the handle, connecting rods, discs and the curved siphon pipe. Components of the latter are the floating ball, the two levers it is connected to, the inlet valve and the water inlet pipe. Turning the handle clockwise creates a chain of events, i.e. component behaviours that causally influence one another, resulting in water being flushed out of the tank through the siphon pipe. This, in turn, causes the water level in the tank to lower, pulling down the floating ball with it. This event generates another causal chain of events resulting in the inlet valve being pulled open and new water entering the tank. These two event chains are distributed in space (i.e. they occur at various locations within the 3D space of the tank) as well as in time (i.e. there is temporal order and concurrency). Understanding these two event chains in the two subsystems and their interaction is crucial to success in comprehending and making predictions about the operation of this device.

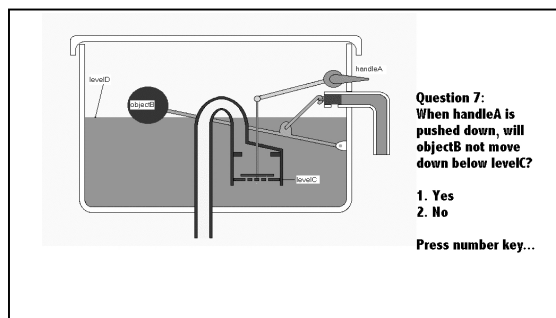


Figure 1. A visuo-spatial and causal reasoning problem.

3 Reactive Information Displays

Interactive displays are increasingly being used for information-oriented tasks in visuo-spatial and causal domains such as emergency response coordination and military strategy planning. But research on displays that attempt to *actively* assist the problem solver in these tasks by determining and responding to one's information needs is still in its infancy. An Attentive User Interface is one that tracks and responds to non-verbal commands and cues from the user. The recent emergence of this research

paradigm attests to the fact that user interface research is moving from a passive, direct-manipulation paradigm toward active and autonomous modes. Attentive User Interfaces can be particularly useful for tasks from visuo-spatial and causal domains. The reason is that the structural and behavioural characteristics of systems in such domains require a systematic problem solving strategy on the part of users who may be trying to understand, predict or troubleshoot these systems. Specifically, the user has to pay attention to the subsystems and components that participate in behaviours and functions of interest, *and* be able to comprehend and predict interacting, spatially and temporally distributed causal event chains that occur in the operation of the system. This presents an opportunity that Attentive User Interfaces which track and respond to attentional cues of the user can profitably exploit.

In this context, we have been researching a subclass of Attentive User Interfaces called Reactive Information Displays. These are displays that leverage four kinds of knowledge in order to provide the right information in the right place and at the right time: (1) knowledge about the system/domain that is being displayed; (2) knowledge about the problem solving task that the user is engaged in; (3) knowledge regarding an applicable problem solving model; and (4) knowledge about the trajectory of the user's attention shifts. For example, such a display for the problem shown in Figure 1 might highlight components participating in the two causal chains of events in the correct order to guide the user's visual attention, and provide additional information (in the form of local animations, for instance) as the user attends to individual components. Reactive Information Displays may be viewed as a kind of non-command [Jacob 1993] or attentive [Vertegaal 2002] information interface.

One of the earliest such displays is described by Starker and Bolt [1990]. Their storytelling display continuously computed a measure of interest for each display object based on the number of user glances. The objects with the highest levels of interest reacted with narrated stories. A cognitive basis for the design of Reactive Information Displays and an empirical study of four kinds of reactive strategies are described in [Narayanan and Yoon 2003]. Experimental results reported in this paper suggest that such displays can aid systematic search and reduce working memory load, thereby enhancing the user's problem solving performance. In particular, the authors found that a display that guides the user's visual attention along paths of causal propagation while displaying potential behaviors of individual components significantly improved the accuracy of mechanical problem solving.

4 Mental Imagery and Animation

Mental imagery is the phenomenon in which someone imagines an object or a visual scene in his or her "mind's eye" in order to retrieve information from that mental image or to transform it so as to generate needed information [Shepard and Cooper 1986]. The more specific process of imagining the behaviors of components of a mechanical system has been termed *mental animation* [Hegarty 1992] and *imagistic simulation* [Clement 1994]. The imagined behavior might be real (e.g. rotating gears) or metaphorical (e.g. flow of electricity). During mental animation, people begin their reasoning with the behaviors of one or a few components, and incrementally infer how these will affect (trigger, modify or prevent) the behaviors of each of the other components of the system. This mental reasoning process is mediated by knowledge about the spatial relations between components, causal relationships among components and domain-specific conceptual knowledge. Empirical evidence suggests that

mental animation is an analog imagery process [Hegarty and Sims 1994; Schwartz and Black 1996].

Narayanan and Hegarty later developed a cognitive process model of how people comprehend multimodal information displays of systems in visuo-spatial and causal domains [Narayanan and Hegarty 1998; 2002], in which mental animation is a crucial process. This model views comprehension as a constructive process with six stages, through which the user integrates his or her prior knowledge of the domain with the externally presented information (on the display) to build a mental model of the system. The resulting internal representation is a mental model that is “runnable” in that it contains information that allows the user to mentally simulate the system and generate predictions about its operation.

5 Investigating Imagery

Narayanan and Hegarty’s model postulates that people construct a mental model of a dynamic system by decomposing it into simpler components, retrieving relevant background knowledge about these components, and mentally encoding the relations (spatial and semantic) between components to construct a static mental model. In the course of solving a specific problem, they “run” this mental model, beginning with some initial conditions and inferring the behaviours of components one by one along causal event chains. This cognitive model influenced our research on Reactive Information Displays because it suggests that mental animation will occur when people are solving a reasoning problem from a visuo-spatial and causal domain that is presented to them on a display. Based on this model, it is possible to generate plausible predictions about the manner or order in which a problem solver is likely to allocate his or her attention to various elements of the information display while solving the problem.

However, for a reactive display to be effective, it has to be able to infer where the user is in the problem solving process at any given moment. In other words, it needs to be able to determine which elements of the display will be attended to next. It is not easy to discern the problem solving focus of a user, i.e. which display element the user is thinking about at any moment. So the question naturally arises as to whether any externally observable manifestations of the user’s behavior can provide clues regarding the shifts in problem solving focus or cognitive attention of a user.

Movement of the eye is one such manifestation that is a likely indicator of what the user is (or has been) thinking about. The assumption (called the eye-mind assumption) is that the locus of eye fixations corresponds to the information being processed by the cognitive system. Just and Carpenter [1976] discuss evidence supporting this assumption for goal-directed tasks that require information to be encoded and processed from the visual environment. Further evidence that eye movement patterns do provide information about cognitive processes that underlie mechanical reasoning appears in [Rozenblit et al. 1998].

Indeed, our earlier work on Reactive Information Displays does suggest that a display that tracks and reacts to eye movements appropriately can enhance the accuracy of problem solving in mechanical reasoning tasks [Narayanan and Yoon 2003]. However, since psychological research makes a distinction between so-called “visualizers” and “verbalizers”, an interesting issue is whether there are users who may benefit more from displays that react to their eye movements. This issue was not considered in our prior research, and it is the focus of the experiment reported in this paper.

In particular, the eye-mind assumption suggests that users’ eye movements will indicate their trajectory of reasoning when solving a problem posed graphically on an information display. This raises the interesting question of whether similar eye movement patterns occur when, after solving such a problem, they are asked to solve a similar problem on the same system but presented only with a blank display. That is, when people have to rely on their memory of a display to solve a causal reasoning problem, will their eye movement patterns provide evidence of mental animation? If eye movement patterns similar to the original display are detected on the blank display also for a particular user, it provides evidence of mental imagery processes at work, and suggests that a display that tracks and reacts to eye movements to provide additional information is likely to help this user in the problem solving process.

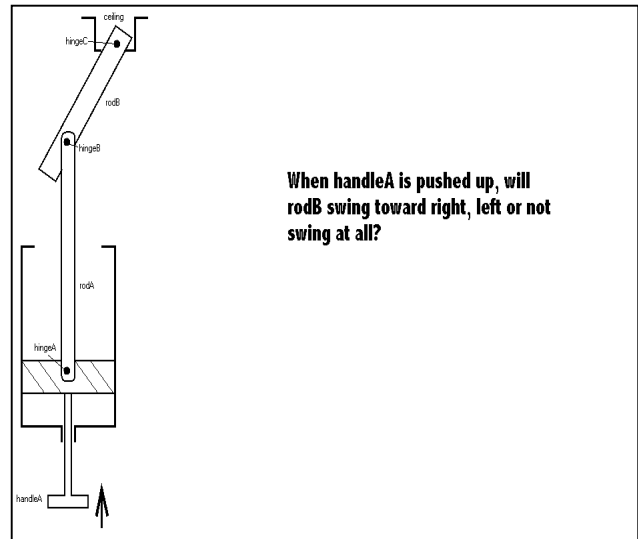


Figure 2. Stimulus display for the first problem.

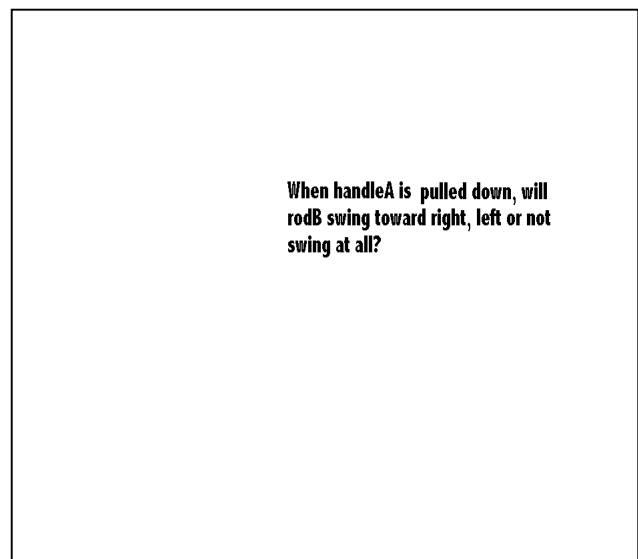


Figure 3. Stimulus display for the second problem.

6 Experimental Procedure

We conducted an experiment to investigate this question. Ninety engineering graduate students volunteered to participate. They were compensated with a payment of \$10 each. The experiment was conducted one subject at a time in an eye tracking laboratory equipped with a head-mounted eye tracker, eye tracking computer and a stimulus display computer. The eye tracker we used is the Eye Link model from SMI Inc. It consists of a headband, to which two infrared sources and cameras (one for each eye) are attached. It is a video-based eye tracker that detects pupil and corneal reflections from infrared illumination to compute screen coordinates of the user's gaze point on the stimulus display monitor once every 4 milliseconds. The headband is attached by cable to a PC which functions as an experiment control station as well as carries out the necessary computations. This PC communicates with the stimulus display computer via an Ethernet link. Subjects sat in a high-backed chair, and viewed the problem on a 20-inch wall-mounted monitor at eye level, at a distance of approximately 3 feet. The experimenter sat behind the subject and controlled the experiment through the eye-tracking computer.

Subjects were asked to solve two problems. In the first problem (Figure 2), there is a picture of a mechanical device and a question. The second problem is a modified version of the first problem. But its display (Figure 3) contains only a question, so subjects have to use their memory of the device. First, subjects read the experimental instruction provided to them on a sheet of paper. This was followed by calibration of the eye tracker. The actual experiment began by the subject clicking the left mouse button to display the first problem. When ready to make the prediction, the subject clicked the left mouse button and then verbally reported the answer to the experimenter. To continue to the next stimulus, subjects clicked the left mouse button again. The mouse clicking allowed the recording of response times. Eye movement data was collected and recorded for both problems and all subjects.

7 Process and Outcome Measures

We collected two process measures (eye movements and response times) and one outcome measure (accuracy of prediction). Response time and accuracy are commonly used metrics of problem solving performance. However, not all problems in visuo-spatial and causal domains have answers that can be unequivocally classified as correct or incorrect. A case in point is developing an action plan for an emergency evacuation from an information display that shows factors such as population distribution, layout of roads, features of the terrain and weather conditions. Here it is as important to ensure that the problem solver has considered all critical elements of the domain as it is to create a feasible plan. Therefore, we developed two derived measures called *coverage* and *order*, in addition to accuracy and response time, to characterize the quality of problem solving. Coverage and order are derived from eye movement data, as explained below.

Coverage is defined as the percentage of objects in the display that were attended to for more than a time interval threshold. We set the threshold to 200 milliseconds, approximately equal to two fixations. Coverage is therefore a number between 0 and 100.

A good strategist will not only attend to all relevant objects in the display, but also consider them in the order that best supports reasoning. For example, a crucial feature that separates expert and

novice problem solving in meteorological reasoning from weather maps is that novices attend to objects that are perceptually salient whereas experts attend to objects that are thematically relevant [Lowe 1999]. Therefore, we developed a metric called order that measures how systematically a user attended to causally related elements of the display. This metric is explained next.

Let S be an ordered sequence of display objects that a user attended to during a problem solving session. So S begins with the first display item attended to, and ends with the last item attended to before the solution to the problem is produced. This sequence is generated from eye movement data. In this sequence, if object j appears immediately after object i , and if i can causally influence j in the event chains of the system, then $i-j$ represents a causal pair in the sequence S . Consecutive causal pairs represent causal subsequences of S . The length of a causal subsequence is the number of causal pairs in it. Order of S is defined as the sum of squares of the lengths of causal subsequences in S . This captures the correctness of the sequential order in which the user visually scanned the display (i.e. each causal pair indicates that the problem solver considered one accurate cause-effect pair of display objects) weighted by the number of consecutive causal pairs that have been considered (i.e. if subjects A and B both considered the same number of causal pairs, but if A looked at longer causal subsequences than B, the value of order will be higher for A than B). Order is a number greater than or equal to zero.

From the raw eye movement data we also computed the total fixation duration on each component using a bounding box technique. One of these components (rodB, the topmost rod connecting the device to the ceiling in Figure 2) is critical to correctly solving the two problems. We additionally determined the total number of fixations of each subject for each problem (excluding fixations on the question and on blank areas of the screen).

8 Results

We found that some subjects exhibited eye movement patterns in the blank display area of Problem 2 similar to their eye movement patterns in Problem 1 (Figure 4 shows an example), while others looked only at the question area in the second problem (Figure 5 shows an example). We categorized the former type of subject behavior as an imagery strategy. We failed to collect eye movement data from 2 subjects due to system problems. Two independent raters looked at the eye movement traces (scan paths) of the remaining 88 subjects for Problem 2 superimposed (bitmap superimposition) on the stimulus for Problem 1 (note that the Problem 2 stimulus was derived from the Problem 1 stimulus by changing two words in the question and blanking out the device diagram). Each rater assigned each subject to an imagery group (i.e. using imagery strategy) or non-imagery group (i.e. not using imagery strategy). The inter-rater agreement was 85%. They disagreed on 13 out of the 88 subjects. For each of the subjects that the two raters disagreed on, the experimenter looked at the subject's scan paths for the second problem superimposed on the Problem 1 stimulus. He then assigned the subject to the imagery group if there was at least one fixation directly on a device component or multiple fixations within the general area occupied by the device diagram. Otherwise (i.e. if there was only one or no fixation within the general area occupied by the device diagram *and* no fixation directly on a component) the subject was assigned to the non-imagery group.

We then compared these two groups' problem solving performance in terms of response time, accuracy, coverage and order. There was no significant difference between imagery group and non-imagery group in terms of accuracy in the first problem ($\chi^2 > 0.1714$, $p = 0.6789$). There was no significant difference between imagery group and non-imagery group in terms of accuracy in the second problem either ($\chi^2 > 0.0094$, $p = 0.9225$).

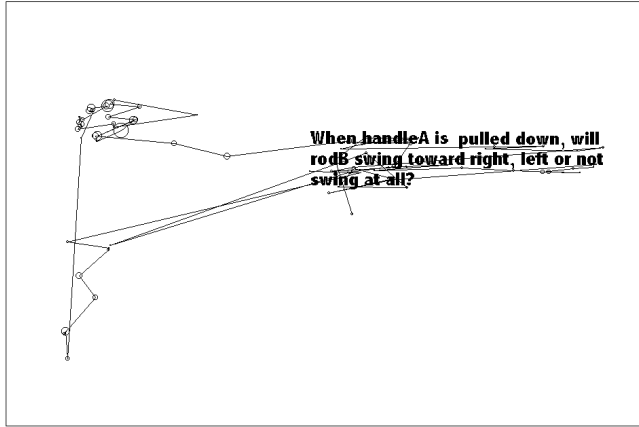


Figure 4. Eye movements indicating an imagery strategy. Subject systematically scanned the empty area occupied by the diagram in Problem 1. The sizes of circles indicate durations of fixations. If the device diagram is re-inserted, it can be seen that many fixations lie directly on or close to components.

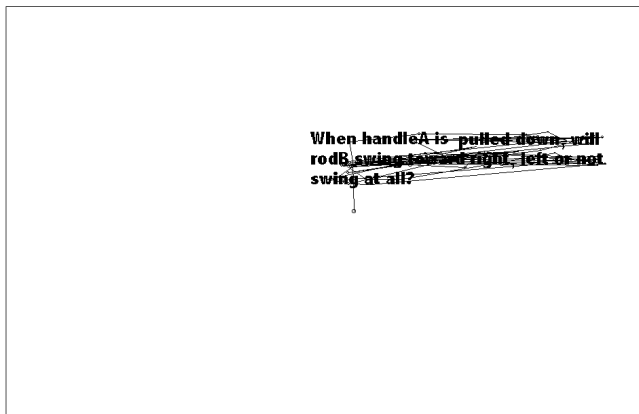


Figure 5. Eye movements indicating a non-imagery strategy. Subject looked only at the question area.

Problem 1	Correct	Incorrect	Total
Imagery	24	13	37
Non-imagery	30	21	51

Table 1. Problem 1 Accuracy

Problem 2	Correct	Incorrect	Total
Imagery	21	16	37
Non-imagery	29	22	51

Table 2. Problem 2 Accuracy

Both Problems	Mean Response Time (secs)	Mean Coverage (%)	Mean Fixation Time on Critical Component (secs)	Mean Value of Order	Mean Number of Fixations
Imagery	38.6	44.6	4.9	74.2	51
Non-imagery	27.3	33.2	3.1	41.9	33.9

Table 3. Dependent Measures

We found significant differences in comparisons of means using the t-test between imagery group and non-imagery group when data from Problem 1 and Problem 2 are considered (Table 3). Imagery group took more time to solve the problem than non-imagery group ($p = 0.0001$). Imagery group had more coverage than non-imagery group ($p = 0.055$). Imagery group also had higher mean fixation duration on the critical component ($p = 0.0386$). Furthermore, the imagery group exhibited higher mean values of order and total number of fixations ($p = 0.0417$ and 0.0364 respectively). These results were replicated with statistical significance when data from only Problem 1 was considered. The imagery group had a higher response time ($p = 0.0375$), more component coverage ($p = 0.0187$), spent more time looking at the critical component ($p = 0.0175$), a higher value of order ($p = 0.0236$), and more number of fixations ($p = 0.0166$).

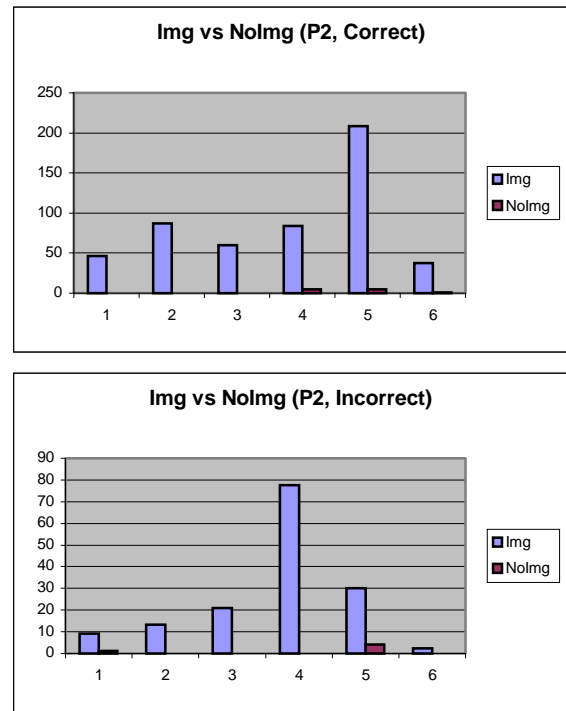


Figure 6. Results of component analysis. Comparison of imagery group (Img) with correct answers to non-imagery group (NoImg) with correct answers in Problem 2 (top). Comparison of imagery group to non-imagery group with incorrect answers in Problem 2 (bottom). X-axis shows each component. The critical component is 5. Y-axis shows time in terms of the number of gaze point data samples (i.e., time in milliseconds = 4 times the number of samples on the y axis).

These results indicate that subjects using the imagery strategy took longer to solve the problem, looked at more components, spent more time looking at the critical component (rodB), and in general visually scanned the display in a more systematic and predictable fashion (i.e. their fixation sequences contained more and longer subsequences of causally related component pairs). Furthermore, these characteristics appeared in Problem 1 (for which the display showed the device diagram) and *persisted* for Problem 2 (for which the display was blank). Thus, imagery subjects exhibited a more systematic visual search of the display, regardless of whether a diagram was present or not during problem solving. This increased their response times, but did not affect accuracy.

In comparing successful subjects to unsuccessful subjects in Problem 2 (imagery subjects only), we found two marginally significant differences in coverage (successful subjects had a higher value of coverage, $p = 0.0829$) and mean fixation duration on the critical component (successful subjects spent more time fixating on the critical component, $p = 0.0707$). These results indicate that imagery subjects who looked at more components and spent more time looking at the critical component tended to be more accurate. This latter conclusion received additional support from the analysis of time that subjects spent looking at individual components (Figure 6). Imagery subjects who were correct in Problem 2 spent the longest time looking at the critical component. However imagery subjects who were wrong in Problem 2 spent the longest time on component 4, which is not a critical component. This implies that correct identification and attention to critical components of a problem can improve accuracy.

9 Concluding Discussion

In the context of increasing research interest in Attentive User Interfaces, and of our own research on Reactive Information Displays, we conducted a study to determine whether there are users who engage in mental imagery during causal problem solving, and whose eye movements reflect their problem solving strategies. In a group of 88 volunteer subjects drawn from engineering graduate students at a public university, we found that 42% can be categorized as using an imagery strategy, as indicated by their scan paths on a region of a blank display upon which they had previously seen a diagram.

While this group did not have an edge in accuracy over those not using an imagery strategy, they took longer to solve the problem, looked at more display elements, and looked at the display for more time in general (as indicated by their total number of fixations). Interestingly, the eye movements of these “visualizers” also indicated a more systematic and predictable strategy of visual attention allocation across elements of the display. That is, they looked at more causally related component pairs and along longer causal chains. Furthermore, the successful problem solvers in the imagery group correctly identified and spent more time attending to a component critical to solving the problems that were posed.

These results have two implications for the design of Attentive User Interfaces in general, and Reactive Information Displays in particular. First, such displays may be of more benefit to users who tend to use the imagery strategy in problem solving than those who do not. Second, reactive displays that help such users better guide their visual attention and reduce the cognitive load inherent in mental imagery by providing additional information locally may improve their problem solving performance (i.e. reduce response time and increase accuracy). For example, a

reactive display might, in addition to showing the image of the device, highlight causally related chains of components in that image, and ensure that the user’s visual attention is guided along these chains by tracking eye movements. It might similarly draw the user’s attention to the critical components or bottlenecks in the problem if it senses that the user did not pay sufficient attention to those. It might provide additional information about the potential behaviors of components, perhaps in the form of local animations that are triggered by fixations, to reduce the user’s cognitive load of mental animation. Experimentally exploring these sorts of reactive strategies is one aspect of our future research.

This experiment has raised several issues in need of further investigation. One such issue is that the device involved in the two problems was very simple. Therefore it is possible that some subjects were able to derive the answer to Problem 2 directly from the understanding they gained in Problem 1 (a learning effect), without resorting to mental imagery or reasoning along the causal chain of events in Problem 2, even though they might have used exactly this strategy in Problem 1. If so, these subjects would not have exhibited eye movements on the blank area in Problem 2, and so would have been classified as non-imagery subjects. But these subjects might very well have used the imagery strategy in Problem 1. Thus the actual percentage of subjects who employed an imagery strategy might be higher than the 42% that we measured. This can be tested by repeating the experiment with a much more complex stimulus. Also, analyses of scan paths with complex stimuli are needed to discover ways of automatically inferring the cognitive state of a user based on his or her trajectory of visual attention across the display.

Other issues are the categorization of subjects as using (or not using) the imagery strategy, and its relation to the classic visualizer/verbalizer distinction in psychology. While we classified problem solvers as using the imagery strategy based only on the location of their fixations on the blank display relative to the device diagram, a finer distinction can be made based on the correspondence between their scan paths and the structure of the device they are reasoning about. We did not pre-classify subjects along the visualizer/verbalizer cognitive style dimension using a questionnaire since we wanted to categorize subjects based on their actual use of mental imagery, as evidenced by eye movements, during problem solving rather than a personality trait. Nevertheless, a match between such a pre-classification and the imagery/non-imagery categorization can provide external validity to the claim that some subjects use mental imagery in problem solving in visuo-spatial and causal domains. These issues will inform our future research as well.

Acknowledgments. Research reported in this paper and the second author’s research program on Reactive Information Displays are supported by the Office of Naval Research under contract N00014-03-10324. We are grateful to the anonymous referees for their constructive comments, which helped improve this paper.

References

- CACM. 2003. Attentive User Interfaces. R. Vertegaal, Ed., *Special Issue of the Communications of the ACM*, 46, 3, 30-72.
- CLEMENT, J. 1994. Imagistic Simulation and Physical Intuition in Expert Problem Solving. In *Proceedings of the 16th Annual Conference of the Cognitive Science Society*, Lawrence Erlbaum, 201-206.

- HEGARTY, M. 1992. Mental Animation: Inferring Motion from Static Diagrams of Mechanical Systems. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 18, 5, 1084-1102.
- HEGARTY, M., AND SIMS, V. K. 1994. Individual Differences in Mental Animation During Mechanical Reasoning. *Memory & Cognition*, 22, 411-430.
- JACOB, R. J. K. 1993. Eye-movement-based Human-computer Interaction Techniques: Toward non-command interfaces. In *Advances in Human-Computer Interaction 4*, Ablex, 151-190.
- JUST, M. A., AND CARPENTER, P. A. 1976. Eye Fixations and Cognitive Processes. *Cognitive Psychology*, 8, 441-480.
- LOWE, R. K. 1999. Extracting Information from an Animation During Complex Visual Learning. *European Journal of the Psychology of Education*, 14, 225-244.
- NARAYANAN, N. H., AND CHANDRASEKARAN, B. 1991. Reasoning Visually about Spatial Interactions. In *Proceedings of the 12th International Joint Conference on Artificial Intelligence*, Morgan Kaufmann, 360-365.
- NARAYANAN, N. H., AND HEGARTY, M. 1998. On Designing Comprehensible Interactive Hypermedia Manuals. *International Journal of Human-Computer Studies*, 48, 267-301.
- NARAYANAN, N. H., AND HEGARTY, M. 2002. Multimedia Design for Communication of Dynamic Information. *International Journal of Human-Computer Studies*, 57, 279-315.
- NARAYANAN, N. H., AND YOON, D. 2003. Reactive Information Displays. In *Proceedings of INTERACT 2003: Ninth IFIP TC 13 International Conference on Human-Computer Interaction*, IOS Press, 244-251.
- ROZENBLIT, L., SPIVEY, M., AND WOJSLAWOWICZ, J. 1998. Mechanical Reasoning about Gear-and-belt Diagrams: Do Eye-movements Predict Performance? In *Proceedings of Mind III: The Annual Conference of the Cognitive Science Society of Ireland*, 158-165.
- SCHWARTZ, D. L., AND BLACK, J. B. 1996. Analog Imagery in Mental Model Reasoning: Depictive Models. *Cognitive Psychology*, 30, 154-219.
- SHEPARD, R. N., AND COOPER, L. A. 1986. *Mental Images and Their Transformations*. Cambridge, MA: The MIT Press.
- STARKER, I., AND BOLT, R. A. 1990. A Gaze-responsive Self-disclosing Display. In *Proceedings of CHI 90*, ACM, 3-9.
- VERTEGAAL, R. 2002. Designing Attentive Interfaces. In *Proceedings of the Eye Tracking Research & Applications Symposium*, ACM, 23-30.