To appear in the *Proceedings of INTERACT 2003: Ninth IFIP TC 13 Int'l. Conf. On Human-Computer Interaction*

Reactive Information Displays

N. Hari Narayanan & Daesub Yoon

Intelligent and Interactive Systems Laboratory, Department of Computer Science & Software Engineering, Auburn University, Auburn, AL 36849, USA

{narayan,yoondae}@eng.auburn.edu

Abstract: Reactive Information Displays (RIDs) that track and react to a user's attention, to present the right information at the right place and in the right time, constitute a relatively unexplored frontier in Human-Computer Interaction. While research on this topic has a two-decade history, advances have been sporadic. One reason may be the difficulty in developing models for predicting the user's attention trajectory and information needs. This paper describes a cognitive model of comprehension and problem solving in visuo-spatial and causal reasoning tasks, points out how principles derived from the model suggest an approach to RID design, and discusses results from an experiment that compared problem solving from a static display and four kinds of RIDs. Results indicate that RIDs have the potential to enhance user performance.

Keywords: attentive interfaces, cognitive model, empirical study, eye tracking, reactive displays

1 Introduction

Increasingly, user interface research is moving from a passive, direct-manipulation paradigm toward active and autonomous modes. One particular strand of research addresses the question of how an interface can *actively* assist the user in information processing tasks. There is a class of problem solving and decision-making tasks for which interactive displays are increasingly being used, but for which research on interfaces that actively assist the user is still in its infancy. This is the kind of task in which 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 accomplish a task such as explaining, planning, predicting or troubleshooting. An example is understanding the weather conditions, and then making a forecast, from an interactive display with various kinds of weather maps, satellite imagery and other information. In these types of tasks, an active interface must leverage knowledge about the task that the user is engaged in and the trajectory of the user's attention shifts in order to provide the right information in the right place and at the right time.

We term such interfaces Reactive Information Displays (RIDs). One of the earliest reactive displays was reported 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 it received. The objects with the highest levels of interest reacted with narrated stories. Such displays may be viewed as a type of non-command (Jacob, 1993) or attentive (Vertegaal, 2002) interface. A theoretical basis for their design and an empirical study of four kinds of reactivity are the foci of this paper. We first present a cognitively motivated theory for RIDs. Then four kinds of reactive displays that we designed to help users solve a prediction problem are described. Finally, an experiment in which we developed two new performance measures and compared user performance from these four displays to that from a static display is reported.

2 Theoretical Framework

2.1 Visuo-Spatial and Causal Domains

RIDs can be particularly useful in domains with tasks that require visual, spatial and causal reasoning (henceforth we will call these VSC domains). These domains share five characteristics: (1) objects of the domain are spatially distributed; (2) the domain is dynamic, i.e. objects and their properties change over time; (3) objects causally interact with each other; (4) such interactions can be traced along chains of cause-effect relationships that branch and merge in spatial and temporal dimensions; and (5) predicting the future evolution of a system in the domain requires reasoning from a given set of initial conditions and inferring these causal chains of events. Examples of domains satisfying these criteria include mechanics, meteorology and military planning.

2.2 A Cognitive Comprehension Model

We developed a model of how people comprehend multimodal information displays of systems in VSC domains. This model, although described in detail with empirical support elsewhere (Narayanan & Hegarty, 2002), is reiterated in the following paragraphs since it provides a theoretical foundation for RID design.

The 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 presented information to build a mental model of the system that is being described. The resulting internal representation is a mental model that is "runnable" in that it contains information that allows the user to mentally simulate a system and generate predictions about its operation.

This cognitive 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. They then mentally simulate this static mental model, beginning with some initial conditions and inferring the behaviours of components one by one in order of the chain of causality or logic. This mental process depends on prior knowledge (e.g. rules that govern the behaviour of the system in question), text comprehension skills, spatial visualization skills for encoding and inferring information from graphic displays, and ability to integrate information in multiple modalities.

Mental model construction under these circumstances involves several stages of processing that are described next. Although these processes are listed sequentially, they are not necessarily accomplished in this order and during comprehension a user may iterate through these processes multiple times to elaborate his or her mental model of the system.

Decomposition. A complex system typically consists of individual components or elements. The user needs to parse the representations on the external display into units that correspond to meaningful elements of the domain. This decomposition process is probably guided both by prior knowledge about the system and its components (a top-down influence) and by visual properties of the external representations.

Constructing a static mental model by making representational connections. Another stage in comprehension involves making representational connections among the visual and symbolic units identified during decomposition. This stage involves making two types of connections: (a) connections to prior knowledge and (b) connections between the representations of different domain elements.

The viewer must make connections between the visual and symbolic elements identified in the external display and their referents. These referents are objects of the domain, as in the case when a blue ribbon on a map refers to a river. This process will be partially based on prior knowledge of both elements of the domain and the conventions used to portray the elements visually or symbolically. The user must also represent the relations between different elements of the system. This involves understanding the physical and causal interactions among components.

Constructing a static mental model by making referential connections. When information is presented in multiple representations and modalities, the user must integrate these different sources to construct a mental model of the system being described. An important process in this integration is resolving co-references, i.e. making referential links between elements in the different representations that refer to the same entity. For example, a military strategist might need to mentally relate a noun phrase "battalion-1", an icon on a map, and a visible troop formation in an aerial photo during planning.

Applying knowledge about basic laws. Complex systems are governed by laws of natural sciences. Some users, especially novices, may lack an understanding of these underlying laws. However, such an understanding is critical for accurate comprehension, especially in the following stages.

Hypothesizing causality. Another stage of comprehension involves identifying the sequence of events, related in terms of cause and effect, which occur in the operation of the system. Our previous studies (Narayanan, Suwa & Motoda, 1995) have shown that people tend to infer events in the operation of a system along directions of causal propagation (lines of action). Finding the causal chains requires knowledge of both the spatial structure of the system and the temporal duration and ordering of events in its operation.

Constructing a dynamic mental model by mental animation and inference. The final stage of comprehension is that of constructing a dynamic mental model by inferring the dynamic behaviours of individual components of the system, and integrating this information to understand how the components work together. Prior work in the mechanical domain (Hegarty, 1992) suggests that this is often accomplished by considering components individually, inferring their behaviours due to influences from other connected or causally related components, and then inferring how these behaviours will in turn affect other components. This incremental reasoning process causes the static mental model constructed in earlier stages of comprehension to be progressively transformed into a dynamic one. This stage can involve both rule-based inferences that utilize prior conceptual knowledge and visualization processes for mentally simulating component behaviours.

2.3 Display Design Principles

The elaboration and application of this cognitive model to information display design in two domains – mechanical systems and computer algorithms – led to six principles for effective information presentations (Narayanan & Hegarty, 2002). These can be described as follows.

Decomposition principle: Provide cues in verbal and visual representations that help users decompose the system or process being explained into simpler components.

Prior knowledge principle: Use words, pictures, sounds, etc. that help users invoke and connect their prior knowledge to the external representations.

Co-reference principle: Use interactive and deictic devices to indicate the common referent if the presentation contains multiple entities in different media with the same referent.

Basic laws principle: When the operation of a system depends on basic principles that might not be understood by all users, describe these principles explicitly in the context of the system being explained.

Lines of action principle: Help a user understand the physical, causal and logical relationships that determine how the behaviour of each part of the system influences that of others.

Mental simulation principle: Encourage and aid users in mentally simulating the system that is being explained, and provide external representations and interactive facilities to reduce their working memory demands during this process.

2.4 Implications for RIDs

Empirical studies (Hansen, Narayanan & Hegarty, 2002) have demonstrated that *interactive displays* designed according to these principles are more effective in aiding comprehension and learning than displays that did not conform to these principles. But now we turn to the issue of how these principles might guide the design of *reactive displays* that actively and autonomously assist the user in a problem-solving task. This avenue of enquiry leads to six specific research questions, corresponding to the six principles.

Is problem-solving performance enhanced when – (1) The user's visual attention is guided along the lines of action? (2) Local behaviours are animated in regions of the display, after the user's attention has been attracted to these regions? (3) The information display is initially sparse, with detailed information progressively revealed? (4) Relevant background information is displayed in regions of the display while the user attends to those regions? (5) Information regarding governing laws is displayed in regions of the display where such laws are deemed to be applicable once the interface determines that the user is attending to those regions? (6) All objects in the display that refer to the same entity are automatically highlighted when the user attends to any one?

While the cognitive model suggests that it is reasonable to expect discernible improvements with each of these display adaptations, these must be empirically tested and validated (or refuted). Thus, the cognitive model provides a theoretical grounding for an empirical and technical research program on designing and testing various kinds of RIDs, and measuring the extent to which such displays improve human problem solving in VSC domains.

The design of RIDs poses two technical challenges. The first involves incorporating knowledge of the tasks and domain in the system to enable it to compute paths of causal propagation, illustrate behaviours of domain objects, determine what background information and governing laws are applicable to which domain objects or regions in the display, and identify the referents of all entities in the display. The second challenge is developing a means of tracking and predicting the shifting attention of the user and applying this knowledge to appropriately control display adaptations.

These are non-trivial problems requiring significant effort, so a prudent and systematic approach to developing effective RIDs is to first empirically investigate each of the six questions. The results of these studies should throw light on what kinds of display adaptations are most and least effective. Once this is experimentally determined, one can be more confident in investing resources to solving the two technical challenges and developing an intelligent architecture for RIDs.

Therefore, in the next section we report on an experiment that investigated the first two questions.

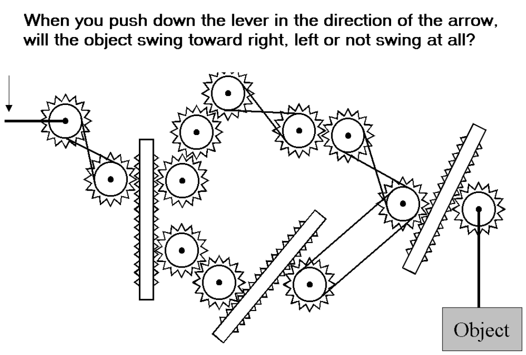

Figure 1: The stimulus.

3 Empirical Study

3.1 Domain and Problem

For this experiment we chose a mechanical reasoning problem, shown in Figure 1. It involves a series of 22 interconnected components: circular and linear gears, belts, levers and a rectangular object. It contains multiple causal propagation paths: starting at the left with the lever, branching into two separate paths at the first linear gear, which then merge into a single path just before the third linear gear, and end at the rectangular object connected to the rightmost lever. The problem is to predict whether and in which direction the object will swing if the leftmost handle is pushed downward, as indicated by the arrow. The correct answer is that it will not swing. Solving this problem requires understanding the different components of the system and their individual (local) behaviours, knowing and applying background information about the different types of components and physical laws that govern their motions, and inferring how pushing the leftmost lever will result in a chain of events that follow the two parallel causal propagation paths and oppose each other at their merge point.

3.2 Four RID Designs

Four kinds of RIDs were built by connecting a head-mounted eye tracker with the 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 (ms). The headband is attached by cable to a computer which functions as an experiment control station as well as carries out computations (such as fixation detection) on eye movement data. This eye-tracking computer communicates with the stimulus display computer via an Ethernet link.

We wrote software running on the stimulus display computer which, in real-time, acquired a user's gaze points from the eye tracking computer and classified these as falling inside predefined bounding boxes of a component, on the text stating the problem to be solved, or on the empty screen area of the stimulus. The four RIDs implemented different approaches to adaptation using this classification.

One, called the *machine-guided highlighting display*, highlighted each component, one after the other, by making it blink. The onset of this behaviour occurred 6 seconds after the stimulus was revealed to a user, to allow him or her enough time to read the problem statement. The blinking effect was achieved by turning the outline of a component red for 400 ms, returning it to the default black colour for 400 ms, and repeating this for 2400 ms for each component. Then, following a delay of 400 ms, the next component would start to blink. Values of the various temporal durations were determined from pilot testing. The goal was to arrive at the smallest possible intervals that seemed sufficient to attract and hold the user's visual attention.

The order in which components blinked followed the paths of causal propagation in the device. Starting with the leftmost lever, the blinking would proceed to the first linear gear at the branching point, then continue along the upper path to the circular gear just before the third linear gear (the merge point). After this object stopped blinking, the first linear gear (the branching point) would start blinking again. Then blinking would continue along the lower path until it reached the rectangular object. After this the whole cycle would repeat. This RID ignored the user's gaze points.

The second, called the *user-guided highlighting display*, employed an identical component blinking behaviour. But it chose the component to blink based on the user's gaze. If the user looked inside the bounding box of a component for more than 200 ms, it would start to blink. If the user then looked at another component for this duration, it would start blinking and the previous one would stop. If the user looked at the question or blank screen space, no blinking would be shown.

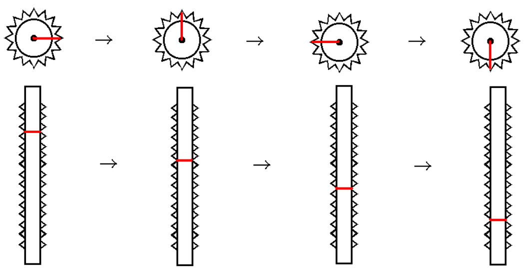

Figure 2: Examples of animation.

The third kind of RID, called a *machine-guided animating display*, showed a localized animation of each component. The onset of this behaviour occurred 6 seconds after the stimulus was revealed to a user, to allow him or her enough time to read the problem statement. The animation effect was achieved for each component by showing a red-coloured line segment in different positions on the component in rapid succession to produce an illusion of motion. A traditional animation (i.e. a component rotating or translating) was not used because we wanted the animations to be localized: when one component was moving, other components it was in contact with did not. So we felt that, given the intermeshed teeth of gears, showing the actual rotation a gear's teeth without it affecting the other gears might produce some cognitive dissonance. So, for a circular gear, a red radial line with an arrowhead would appear at 0, 90, 180 and 270 degrees (each line would persist for 400 ms) to indicate one complete clockwise rotation. The lines would appear in the opposite order for counter-clockwise rotation (Fig. 2). The four lines used to indicate motion of a linear gear can also be seen in Fig. 2. They appeared and disappeared with the same temporal durations, and in an order commensurate with the direction of the linear gear's motion. Similarly, six red line segments (persisting for 400 ms each) were used to animate a crossing belt and 4 segments (persisting for 400 ms each) were used to animate a non-crossing belt. After the last line segment was switched off, the first line segment would appear again. These animation cycles would be repeated for a total of 2400 ms for each component. Then, following a delay of 400 ms, the next component would start to animate. Values of the various temporal durations were determined from pilot testing with the goal of arriving at the smallest possible intervals that seemed sufficient to attract attention and convey a sense of motion.

The order in which components were animated followed the paths of causal propagation in the device. Starting with the lever, the animation would proceed to the first linear gear at the branching point, then continue along the upper path to the circular gear immediately to the left of the linear gear at the merge point. After this gear stopped animating, the first linear gear (the branching point) would start animating again. Then animation would continue along the lower path until it reached the circular gear immediately to the left of the linear gear at the merge point again. After this the whole cycle would repeat. Note that animations of the linear gear at the merge point or of any component thereafter were not shown in order not to give away the solution to the problem. This RID ignored the user's gaze points.

The fourth, called the *user-guided animating display,* employed identical component animating behaviours. But it chose the component to animate based on the user's gaze behaviour. If the user looked inside the bounding box of a component for more than 200 ms, it would start to animate. If the user then looked at another component for this duration, it would start animating and the previous one would stop. If the user looked at the question or blank space, no animation would be shown.

3.3 Experimental Conditions

We measured human problem solving performance on this problem under five experimental conditions using one control display and the four RIDs described above. In the control condition, subjects saw a static display of the stimulus (Fig. 1) on a computer monitor.

The designs of these four RIDs were based on the fact that this experiment was designed to probe two of the questions posed in Section 2.4. The machine-guided highlighting display was intended to investigate the question of whether problem-solving performance can be improved by guiding a user's visual attention along the lines of action (Question 1, Section 2.4). The specific attention-drawing mechanism, component outline blinking, was chosen because prior research on solving a diagrammatically posed medical problem (Karl Duncker's radiation problem) showed that attracting the user's attention to regions important to the problem at hand with a blinking behaviour resulted in improved accuracy (Grant & Spivey, 2002). The regions important to correctly solving the problem in Figure 1 are the contacting surfaces of components, which appear in the diagram as outlines of the components. The control (a static display) and the machine-guided highlighting display differed in two ways. First, the outlines of the components blinked, attracting attention to the contacting surfaces. Second, they blinked along the paths of causal propagation.

Grant and Spivey showed that merely attracting attention to relevant regions can improve performance (the problem they used was a simple one in which causal order did not play a part). Therefore, the user-guided highlighting display was designed as a contrast to the machine-guided highlighting display. That is, to detect the effect of attracting the user's attention to the contacting surface of a component that he or she is already looking at, *without* revealing the lines of action in the device. The static display and the user-guided highlighting display differed only in terms of the absence or presence of the blinking behaviour.

The user-guided animating display was designed to investigate the question of whether problem-solving performance can be improved by animating local behaviours after a user's visual attention is attracted to components (Question 2, Section 2.4). So this display did not reveal the lines of action in the device. Thus, the static display and the user-guided animating display differed only in terms of the absence or presence of component animations.

The machine-guided animating display was intended to investigate the question of whether problem-solving performance can be improved by *both* animating local behaviours and pointing out the lines of action. The static display and the machine-guided animating display thus differed in two ways: in animated illustration of local behaviours and showing the paths of causal propagation.

3.4 Process and Outcome Measures

We collected two raw process measures (eye movements and response time) and one outcome measure (accuracy of prediction). Response time and accuracy are commonly used metrics of problem solving performance. However, not all problems in VSC 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. In this case it is as important to ensure that the problem solver has considered all critical elements of the domain, as it is to create a workable plan. Therefore, we developed two new measures called *coverage* and *order* 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. It is a number between 0 and 100. We set the threshold to 200 ms, since eye fixations are often in the range 200-400 ms.

A good problem solver 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, order is a metric that measures how systematically a user attends to causally related elements of the display.

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 computed from fixations detected by the eye tracker. In this sequence, if object j appears immediately after object i, and if i can causally influence j according to the lines of action in the system that the display shows, then i-j represents a causal link in the sequence S. Consecutive causal links represent causal subsequences of S. The length of a causal subsequence is the number of causal links 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 order in which the user processed objects in the display (i.e. each causal link indicates that the problem solver considered one accurate cause-effect pair of display objects) weighted by the length of lines of action that were considered (i.e. if users A and B both considered the same number of causal links, but if A looked at longer chains of causal links than B, the value of order will be higher for A than B). Order is a number greater than or equal to zero.

3.5 Experimental Procedure

90 graduate students of engineering volunteered to participate in the experiment, in return for a nominal payment. They were recruited through emailed advertisements. Subjects were randomly assigned to one of 5 conditions: static display (S; n=15) machine-guided highlighting display (MH; n=15), user-guided highlighting display (UH; n=20), machine-guided animating display (MA; n=20) and user-guided animating display (UA; n=20). The experiment was conducted one subject at a time in our eye tracking laboratory equipped with a head-mounted eye tracker, eye tracking computer, and a stimulus display computer.

Subjects sat in a high-backed chair, and watched the stimulus display on a 20-inch monitor mounted on a wall 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. First, subjects read a description of the experiment and solved two practice problems. These were mechanical reasoning problems with very simple devices that did not contain any component type that appeared in the experimental stimulus (Fig. 1). A 5-minute break and calibration of the eye tracker followed. The actual experiment began by the subject clicking the left mouse button to display the stimulus. When ready to make the prediction, the subject depressed one of three designated keys on the keyboard to indicate his or her answer. The mouse and key presses allowed the recording of response time and determination of accuracy. Eye movement data was also collected and recorded.

3.6 Results

Table 1 shows mean values of accuracy (A: percentage), response time (RT: seconds), coverage (C: percentage) and order (O: number) for the 5 conditions. Figures 3 and 4 show accuracy and response time charts with error bars. We computed the Chi-square statistic on accuracy differences between the control condition and each of the four RID conditions, and between each pair of the RID conditions. The machine-guided animating display produced a statistically significant improvement in accuracy over the static display (χ^2=7.7038,p=0.0056,n=35). We computed the t-test on differences in mean response time between the control condition and each of the four RID conditions, and between each pair of the RID conditions. The machine-guided animating display induced a marginally significant increase in mean response time compared to the static display (t(34)=1.882,p=0.0687).

	S	MH	UH	MA	UA
A	40%	60%	60%	85%	60%
RT	76	112	100	114	86
C	55.5%	59.1%	63.4%	63.1%	56.4%
O	80.2	83.2	107.8	106.7	89.5

Table 1: Results.

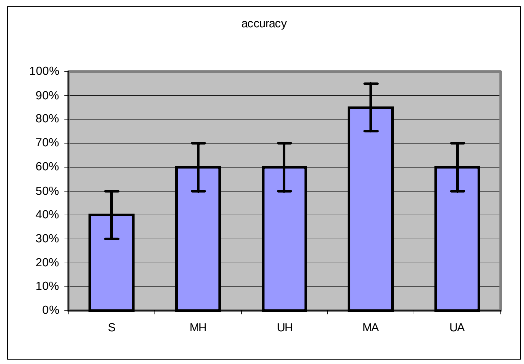

Figure 3: Accuracy for the five conditions.

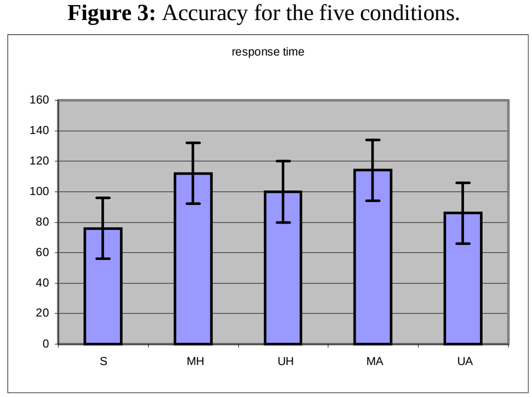

Figure 4: Response Time for the five conditions.

3.7 Discussion

This experiment investigated the performance impacts of RIDs that actively guide the user's attention along paths of causal influence and provide localized information on component behaviours. Based on the theory, we expected that RIDs would produce better accuracy, coverage and order than a static display. We also expected that RIDs would increase response time since these encourage a user to attend to all display objects systematically. As Table 1 and Figs. 3 and 4 show, all RIDs produced better problem solving accuracy and required more response time than the static display. The coverage and order values in the table suggest that RIDs also induced more systematic visual search than the static display. The accuracy for the machine-guided animated display was significantly higher than that for the static display. This suggests that any one of the reactive strategies – highlighting the relevant parts or showing the local behaviour of a display object that a user is already looking at, or pointing out the lines of action in the system – may not by itself be sufficient to produce significant improvements in accuracy. But an RID that shows local behaviours as well as paths of causal propagation (i.e. one that addresses *both* questions 1 and 2, Section 2.4) can significantly improve the problem solving performance of users.

4 Conclusion

This paper makes two contributions. The primary one is a theoretical framework for the principled design and evaluation of information displays that track the visual attention of users and react according to a generalized model of comprehension and problem solving in VSC domains. This, in turn, raised six research questions about the impact of different kinds of reactivity. An experiment that investigated two of these questions is reported. Its results indicate that two kinds of reactivity (showing local behaviours and lines of action) can improve accuracy, by reducing working memory load and encouraging systematic search. This is the second contribution.

Our future research will address the other four questions and evaluate new RID designs in not only mechanical but also other VSC domains. The machine-guided and user-guided displays present two extremes: the former is blind to the user's visual search patterns whereas the latter simply reacts to the user's present focus of attention. Designing and testing RIDs that are more intelligent is an important next step. One example is a display that, based on past gaze behaviour, determines the branch of a causal path the user has been on (or, alternately, predicts the branch that will be of interest next) and then guides the user in traversing that path further.

Acknowledgment: Office of Naval Research supported this research under contract N00014-03-10324.

References

Grant, E. R. & Spivey, M. J. (2002), Guiding attention produces inferences in diagram-based problem solving, *in* M. Hegarty, B. Meyer & N. H. Narayanan (eds.), *Diagrammatic Representation and Inference*, LNAI 2317, Springer, pp.236-248.

Hansen, S. R., Narayanan, N. H. & Hegarty, M. (2002), Designing educationally effective algorithm visualizations, *Journal of Visual Languages and Computing*, 13(3), 291-317.

Hegarty, M. (1992), Mental animation: Inferring motion from static diagrams of mechanical systems, *JEP: Learning, Memory & Cognition*, 18(5), 1084-1102.

Jacob, R. J. K. (1993), Eye-movement-based human-computer interaction techniques: Toward non-command interfaces, *Advances in Human-Computer Interaction*, Vol. 4, Ablex, pp.151-190.

Lowe, R. K. (1999), Extracting information from an animation during complex visual learning, *European Jnl. Psych. of Education*, 14, 225-244.

Narayanan, N. H. & Hegarty, M. (2002), Multimedia design for communication of dynamic information, *Int'l. Jnl. Human-Computer Studies*, 57, 279-315.

Narayanan, N. H., Suwa, M. & Motoda, H. (1995), Diagram-based problem solving: The case of an impossible problem, *Proc. Annual Cognitive Science Society Conference*, Lawrence Erlbaum, pp.206-211.

Starker, I. & Bolt, R. A. (1990), A gaze-responsive self-disclosing display, *Proc. Human Factors in Computing Systems Conference*, ACM, pp.3-9.

Vertegaal, R. (2002), Designing attentive interfaces, *Proc. Eye Tracking Research & Applications Symposium*, ACM, pp.23-30.