Diagrams in the Mind and in the World: Relations between Internal and External Visualizations

Mary Hegarty

University of California, Santa Barbara Santa Barbara, CA 93106
hegarty@psych.ucsb.edu

Abstract. Recent advances in computer technology and graphics have made it possible to produce powerful visualizations of scientific phenomena and more abstract information. There is currently much excitement about the potential of these computer visualizations, particularly in education and training in science, technology and medicine. This paper examines three possible relations that might exist between internal and external visualization. First, external visualizations might substitute for internal visualizations. Second their comprehension and use may depend on internal visualizations. Third, they might augment and be augmented by internal visualizations. By reviewing these possibilities, it is argued that the design of external visualizations should be informed by research on internal visualization skills, and that the development of technologies for external visualizations calls for more research on the nature of internal visualization abilities.

Introduction

Recent advances in computer technology and graphics have made it possible to produce powerful visualizations of scientific phenomena and more abstract information. There is currently much excitement about the potential of these computer visualizations, as indicated by the publication of several recent reports and books on scientific and information visualization (Nielson, Shriver & Rosenblum, 1990; Card, Schneiderman & McKinley, 1999 Spence, 2001). These external representations are seen as having the potential to augment, enhance, or amplify cognition (Card et al, 1999; Norman, 1993) and educational theorists have stressed the need for exposing students to these powerful visualizations in science education.

The goal of this paper is to examine the possible relations that might exist between external visualizations and internal visualizations. An external visualization is an artifact printed on paper or shown on a computer monitor that can be viewed by an individual. An internal visualization is a representation in the mind of an individual. For the purposes of this paper, I define a visualization as any display that represents information in a visual-spatial medium. Although computer visualizations are often dynamic and interactive, my definition also includes static displays such as graphs, diagrams and maps. Internal and external visualizations can represent physical phenomena that are spatial in nature, such as the development of a thunderstorm. They can also depict more abstract phenomena, such as the flow of information in a computer program or the organization of information on the world-wide web.

A. Blackwell et al. (Eds.): Diagrams 2004, LNAI 2980, pp. 1-13, 2004.

Although there is much excitement about the potential of external visualizations at present, the design of external visualizations is often seen as purely a computer science problem, with the goal of designing systems that can present sophisticated and realistic visualizations. Implicit in this view is the idea that external visualizations do not just augment but can perhaps replace internal visualizations. I argue that the potential for developing external visualizations does not replace the need to understand or develop internal visualization skills. On the contrary, I argue that the design of external visualizations should be informed by research on internal visualization skills, and that the technological developments call for more research on the nature of internal visualization abilities.

External Visualizations. In recent history, technology has significantly improved our ability to create external visualizations. It is probably not surprising therefore that current research on the role of visualization in thinking focuses on external visualizations. External visualizations are seen as important ways of augmenting human cognition. This is evident in Card, Schneiderman & McKinley's (1999) definition of visualization as “The use of computer-supported interactive visual representations of data to amplify cognition (p 6)” or Norman's (1993) statement that “the real aids come from devising external aids that enhance cognitive abilities (p. 43)”. Educational research emphasizes the need to expose children to powerful external visualizations of data. For example, in a recent article on the potential of scientific visualization for education, Gordon and Pea (1995) concluded:

> “The case has been made that SciV shows remarkable potential to help students learn science. As an extraordinary plastic medium, SciV affords the construction of provocative images that can resemble physical phenomena and serve as the basis for the construction, debate, and negotiation of meaning that stands at the heart of the process of education” (p 276)

Cognitive scientists have made important contributions to recent research on external visualizations. One is a set of theoretical proposals for why external visualizations are effective. For example, in a seminal article entitled “Why a diagram is worth 10,000 words” Simon and Larkin (1987) argue that diagrams allow insights to be gained by powerful and automatic perceptual inference processes, and that they perceptually group information that must be integrated mentally (see also Koedinger & Anderson, 1990) Another is a specification of principles that make external visualizations more or less effective (e.g., Cheng, 2002; Kosslyn 1994; Shah, Mayer & Hegarty, 1999). Cognitive scientists have also developed models of how humans comprehend and make inferences from external visualizations (e.g., Pinker, 1990; Carpenter & Shah, 1998; Hegarty, 1992).

Internal Visualizations. There has also been an important tradition of research in cognitive science on internal visualization, that is, our ability to internally represent objects, events and more abstract phenomena as mental images, and our ability to infer new information by transforming these images. Studies of internal visualization, have examined people's ability to construct, inspect and transform mental images (Kosslyn, 1980). Shepard & Metzler's (1971) classic research on mental rotation showed that mental transformations are isomorphic to physical transformations of objects. More recent research, using both behavioral methods (Finke, 1989) and cognitive neuroscience methods (Kosslyn, 1994) has demonstrated that imagery and perception involve some shared mechanisms. Our ability to construct and maintain internal visualizations has also been an important topic in the working memory literature (Baddeley 1986; Logie, 1995). In this literature, the internal visualization system is called the visual-spatial sketchpad, and has been shown to be important in such activities as visual imagery, playing chess, and playing video games (Logie, 1995). Finally, the literature on human intelligence has identified spatial ability as one of the most important components of human intelligence. This research suggests that there may be several somewhat dissociable spatial abilities, but the most robust and well-documented spatial ability, known as *spatial visualization ability*, is defined as reflecting "processes of apprehending, encoding, and mentally manipulating spatial forms" (Carroll, 1993, p. 309). Examples of tasks used to test spatial visualization ability are shown in Figure 1. All of these tasks involve imagining the result of spatial transformations, such as rotating an object or folding a piece of paper.

Psychological research on internal visualization has been more concerned with the nature of the imagery system than in its role in thinking and reasoning. However, research in the history and philosophy of science indicates that internal visualizations have been quite powerful in scientific discoveries and inventions (Ferguson, 1977; Miller 1986; Nersessian, 1995). For example, Einstein claimed that he rarely thought in words, and had difficulty translating his images into equations and verbal communications (Miller, 1986). Similarly, ability to mentally simulate the behavior of machines has been reported to be central to design and invention (Ferguson, 1977; Shepard, 1978). Nikolai Tesla, Oliver Evans (inventor of the automatic flour mill) and Walter Chrysler (founder of the automobile company) are among the many famous engineers who claimed to mentally imagine their inventions in great detail before beginning to build them (Ferguson, 1977). Research on individual differences in spatial visualization also points to the importance of internal visualization in science, engineering and medicine. Spatial visualization ability, which can be seen as a measure of internal visualization, is correlated with success in mechanical occupations, mathematics, physics and medical professions (Hegarty & Waller, in press).

Relations between Internal and External Representations. The purpose of this paper is to explore the relationship between internal and external visualizations and the implications of this relationship for education and training in scientific, engineering and medical professions. I consider three different relations that external visualizations might have to internal visualizations. One possibility is that external visualizations can *substitute* for internal visualizations. That is, a person can have the same insight, (or perhaps a better insight) by viewing or manipulating an external visualization of some phenomenon as he or she would have by internally visualizing the same phenomenon. If this is true, an external visualization can act as a "prosthetic". It is not necessary for someone to have internal visualization skills, and education should focus on exposing people to many external visualizations rather than developing the ability to visualize internally. A second possibility is that use of external visualizations *depends* on the ability to internally visualize. In this case, gaining insight from an external visualization would depend on the same abilities and skills as internal visualization. Fostering internal visualization would be an essential goal of education and training. Finally, a third possibility is that external visualizations *augment* internal visualizations, by providing information or insights that are additional to those that can be inferred from internal visualizations. In this case there is continuity between what can be internally visualized and what can be learned from an external visualization, and education should both foster internal visualization and expose students to external visualizations.

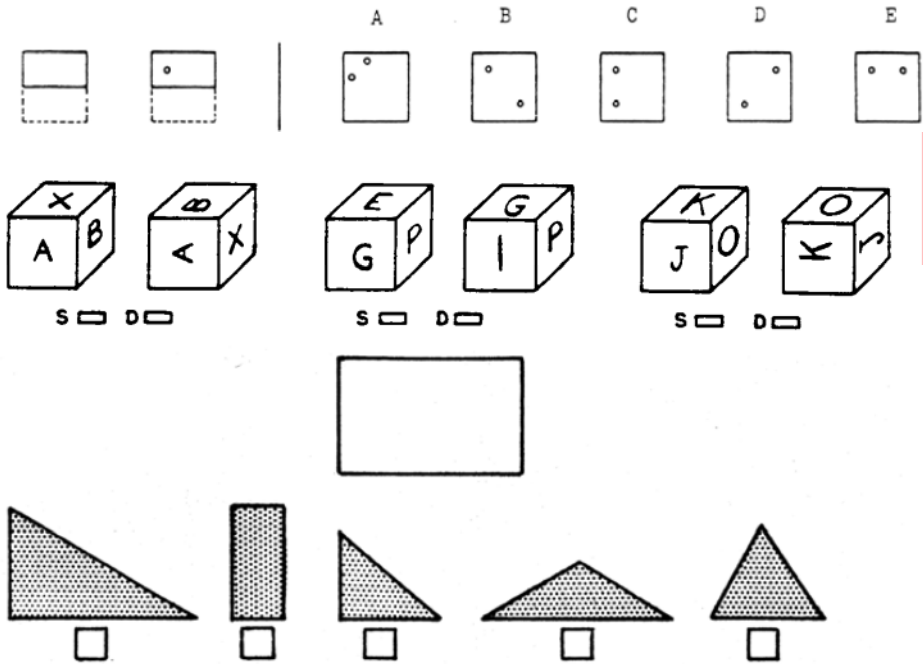

Fig. 1. Examples of items from tests of Spatial Visualization ability In the first test (the Paper Folding Test), the test taker has to imagine that a piece of paper is folded and a hole is punched through it and must choose which of the figures on the right will result when the paper is then unfolded. In the second example (the Cube Comparisons test), the test taker has to decide whether the cube on the right could be a rotation of the cube on the left. In the third example (the Form Board test) the test taker must decide which of the shapes on the bottom would be needed to fill the rectangle on the top, assuming that the shapes can be rotated to fit

These three possibilities are not mutually exclusive. Each of the possibilities might be true for different types of visualizations, different types of content, or different types of people. For example, viewing a data graph may act as a substitute for internally visualizing the graph, but understanding an animation of a thunderstorm may be possible only for those with high spatial visualization ability. Furthermore, external visualizations might both augment internal visualizations and depend on internal visualization ability. However, for the purposes of exposition, I will consider them as three separate hypotheses. I will each consider each of the three situations in more detail and examine where some cognitive studies on visualization fit within this framework. I will draw on research concerning computer animations of mechanical systems, visualizations of anatomy used in medical education and practice, and visualizations of meteorological phenomena. Research on these topics has examined how people comprehend and learn about phenomena that occur in three-dimensional space and change over time by viewing visualizations that are presented in the two dimensional media of the printed page and the computer monitor.

Relation 1: External Visualization is a Substitute for Internal Visualization

The first possibility is that viewing an external visualization of a phenomenon can be a substitute for internally visualizing the same phenomenon. If this is true, the availability of external visualizations relieves us of the necessity of internally visualizing and an external visualization can serve as a "prosthetic" for those who have difficulty with internal visualization, so there is no need to be concerned about the nature of internal visualizations or to develop internal visualization abilities among students. This is perhaps a straw-man theory, but it is implicit in any program of research or practice, which focuses exclusively on the development of external representations as a means of educating and training students without regard to their internal visualization abilities.

Let us explore what it would mean for external visualizations to substitute for internal visualizations in this way. First, it would imply that if people have access to interactive external visualizations, they would manipulate these visualizations and therefore no longer need to carry out effortful internal visualization processes. However, recent research suggests that this is not what happens. Trafton and his colleagues (Bogacz & Trafton, 2002; Trafton, Tricket & Mintz, in press; Tricket & Trafton, 2002) have examined how expert scientists and meteorologists interact with external visualizations in problem solving situations, such as weather forecasting and data analysis. A striking result of these studies is that even in situations in which the experts can manipulate an external display to reveal new data or a new perspective on some phenomenon, they rely extensively on internal visualizations, and in some studies, they manipulate internal visualizations more often than they use the computer interface to manipulate the external display. It is clear therefore that external visualizations do not substitute for internal visualizations for these experts.

It might be argued that external visualizations are more likely to substitute for internal visualizations in the case of novices, who have not yet developed the ability to internally visualize. For example, by viewing an external visualization of some scientific phenomenon, a novice might quickly gain some insight into that phenomenon. This view assumes that if an individual (Person 1) has an insight that is based on or at least accompanied by an internal visualization, and externalizes that visualization, a second individual (Person 2) can view that external visualization and have the same internal representation and insight as Person 1. In the extreme, this would be true, regardless of the abilities or prior knowledge of Person 2. For an external visualization to substitute for an internal visualization in this way, the following conditions would have to hold. First, the external representation of a phenomenon would have to be isomorphic to its internal representation. Second, merely viewing the external representation would have to be sufficient to recreate an internal copy of this representation.

Are Internal Representations Isomorphic to External Representations? The idea that the internal representation of a phenomenon is isomorphic to its external representation is related to “the resemblence fallacy” proposed by Scaife and Rogers (1996). Research on mental representations of motion in simple mechanical systems, demonstrates how difficult it would be to create an external visualization that is isomorphic to a person's internal representation of a mechanical system. When a machine (e.g. a pulley or gear system) is in motion, many of its components move at once. Therefore, most people would agree that a realistic computer visualization representing the motion of objects in a mechanical system would also show its components moving simultaneously. In contrast, when people attempt to internally visualize how a machine works, they infer the motion of components one by one, in order of a causal chain of events (Hegarty, 1992; Hegarty & Sims, 1994, Narayanan, Suwa & Motoda, 1994; 1995). That is, their mental representation is of a causal sequence of events rather than of a set of simultaneous events. In externally representing how a machine works, one therefore has to choose between showing the components moving sequentially (as the process is conceptualized by most people) or simultaneously (as they move in the real world). Furthermore, although inferring the motion of mechanical systems often depends on spatial visualization processes (Hegarty & Sims, 1994; Sims & Hegarty, 1997, Schwartz & Black, 1996) it can also be based more on verbal rules of mechanical inference, so that most current theories of people's mental models of machines allow for hybrid representations involving both visual-spatial representations and verbal rules (Schwartz & Hegarty, 1996; Narayanan, Suwa & Motoda, 1994). Again, it is not clear how to translate such a hybrid representation into an external visualization.

A related problem is that people's internal mental models of spatial objects and events are not purely visual-spatial, but may depend on other senses. For example, surgeons often report that they “see” with their hands as much as with their eyes during surgical procedures, suggesting that their mental representations of anatomy are based on both haptic and visual modalities. As a result, many common errors in minimally invasive surgery, carried out using instruments that are inserted through small incisions in the skin, are attributed to loss of the haptic information that is available in open surgery. Use of pure visualizations, either in teaching anatomy or in training surgical procedures may therefore lead to an impoverished internal representation compared to the experience that it is intended to simulate. These examples indicate that it may not be possible to externally represent the internal representation of a complex phenomenon in the visual modality alone.

Does Viewing of an External Visualization Lead to a Veridical Internal Representation? The existence of visual illusions is evidence that perception of visual stimuli does not necessarily lead to a veridical internal representation of these stimuli. For example, an object in free fall is seen as falling with a constant velocity, although it is in fact constantly accelerating, and an object moving at a constant velocity is not seen as constant (Runeson, 1974). Such misperceptions are not just true for perception of apparent motion, as shown by a rapid sequence of static images in an animation, but also for perception of real motion. For example, before stop-gap photography, nobody knew how a horse's legs move while it is galloping, because they move too fast to be perceived accurately. As a result, the legs of galloping horses have been portrayed incorrectly in art throughout history (Tversky, Morrison, & Betrancourt 2002). Therefore, viewing a motion event, either in the real world or in an animation, does not necessarily lead to an accurate internal representation of the event.

The internal representation that results from viewing an object, scene, or event can also be dependent on the knowledge of the viewer. For example, when a chess grandmaster views a chess board in the middle of a game, he sees a small number of visual chunks representing groups of chess pieces in attack and defence relations to each other (Chase & Simon, 1973). In contrast, a novice chess player viewing the same chess board sees a large number of individual chess pieces. Similarly, perception of visual representations such as graphs, maps and animations can be influenced by knowledge of both the phenomena depicted in those representations and the graphical conventions used to visualize those phenomena (e.g. the use of color to show temperature and isobars to show pressure on a weather map). For example, Lowe (1999) has found that novices interpretations of weather maps and animations are largely descriptions of surface features of the maps and do not contain the types of insights into meteorological phenomena that are evident in experts' interpretations of the same maps and animations. Therefore if the creator of a visualization and the viewer of the same visualization differ in knowledge, it is very unlikely that the viewer will end up with the same mental representation of the referent object or event as that of the creator of the visualization.

In summary, there is no reason to assume that the internal representation of an object or event is isomorphic to the object or event itself or to its external visualization. Similarly, there is no reason to assume that viewing an object or event or its external visualization will lead to a veridical internal representation. This suggests that the development of external visualizations alone is not a solution to how to communicate about complex phenomena. This design must also take account of the nature of people's internal visualizations.

Relation 2: External Visualization as Dependent on Internal Visualizations

Rather than replacing internal visualizations, the use and comprehension of external visualizations might be dependent on internal visualization skills. In this case people with less spatial visualization skill or lower spatial ability would have poorer comprehension of external visualizations than those with more skill or more ability. A recent study on a classic visual illusion is a case in point. Isaak and Just (1995) studied ability to judge the trajectory of a point on a rolling ball, which involves both translation and rotation. In this situation, people are subject to an illusion called the curtate cycloid illusion, which can be explained by a model in which they temporarily fail to process the translation component of motion at a critical point in the rolling motion. Isaak and Just found that people with high spatial visualization ability were less subject to the illusion than people with low spatial visualization ability. They proposed that spatial working memory, necessary for generating internal visualizations, was also necessary to simultaneously process the rotation and translation components of motion in comprehension of the visual display. In complex visualizations in which several motions occur at once or the motion of several objects must be tracked, people with better internal visualization skills appear to have an advantage in comprehension.

Learning from external visualizations may also depend on ability to visualize internally. Mayer and Sims (1994) examined the role of spatial ability in learning from animations that explain how a mechanical system (car brakes) and a biological system (the human respiratory system) work. They considered two alternative hypotheses. The first hypothesis was that viewing an animation would compensate for low spatial ability (i.e., the animation would act as a prosthetic for those with poor internal visualization skills, as discussed above) so that low spatial individuals would learn more from animations than they do from static diagrams and the differences in learning between high- and low spatial individuals would be smaller for animations than for static diagrams. The second hypothesis was that spatial ability would be a necessary prerequisite for learning from an animation. In this case high-spatial individuals would learn more from animations than low-spatial individuals and the difference between high- and low-spatial individuals would be greater in the case of animations than in the case of static diagrams. The results were consistent with the second hypothesis. That is people with high spatial ability learned more from the animation than people with low spatial ability, and the difference between learning of high- and low-spatial individuals was greater in the case of animations.

In the field of medical education, spatial abilities have been found to be related to the ability to learn anatomical structure from visualizations that reveal the three-dimensional structure of the anatomy by showing rotating models. Garg, Norman, Spero & Maheshwari (1999) found that exposing students to non-interactive three-dimensional models of carpal bone configurations impaired spatial understanding of the anatomy for low-spatial students while it enhanced understanding for high-spatial students. This can be seen as another example of a situation in which internal visualization ability is a necessary prerequisite for learning from an external visualization.

Finally, in some recent research in my laboratory, Cohen, Hegarty, Keehner and Montello (2003) examined students' ability to draw the cross-section of an object that results when the object is cut along a particular plane. Initial studies of this ability indicated that it is highly related to spatial ability. In a follow up study, we allowed students to interact with a three-dimensional computer visualization of the object as they drew the cross section. The students were allowed to interact with the computer visualization to rotate the object, so that they could see it from different perspectives. This study again replicated the high correlation between spatial visualization ability and ability to draw the cross sections. A more surprising result is that there was also a significant correlation between spatial ability and use of the visualization such that the low-spatial individuals rotated the visualization less often than the high-spatial individuals. When interviewed, some of the low-spatial individuals indicated that they did not understand how the visualization could help them. This study indicates that in some situations, internal visualization abilities might be a prerequisite for being able to use an external visualization effectively.

Relation 3: External Visualizations as Augmentations of Internal Visualizations

As reviewed above, recent studies of external cognition describe external visualizations as amplifying, enhancing or augmenting cognition (Card, McKinlay & Schneiderman, 1999; Norman, 1993). This view suggests that there is an interplay between internal visualization processes and comprehension of external visualizations, or that insight or learning is based on a combination of internal visualizations and perception of external visualizations. Although there has been important research specifying how internal and external static representations interact in the performance of cognitive tasks (Zhang, 1996; Zhang & Norman, 1994) this work has not focused on visualizations or their roles in learning about spatial structures or dynamic processes.

One possibility is that by externalizing a mental visualization, a person relieves him or herself of the need to keep the internal representation in working memory, and this frees up processing resources for making further inferences (cf. Zhang & Norman, 1994). For example, Hegarty & Steinhoff (1997) studied the ability of high- and low-spatial individuals to "mentally animate" or infer the motion of components in a mechanical system from a static diagram of the system. Previous studies had shown that low-spatial individuals were poorer at mental animation and we suggested that this was because they had less spatial working memory and therefore could not maintain a spatial representation of the machine while simultaneously imagining the motion of its components. Hegarty and Steinhoff examined the effects of allowing high- and low-spatial individuals to make notes on diagram of mechanical systems while they inferred the motion of components. They found that low-spatial individuals who made notes indicating the direction of motion of each component in the mechanical system were more successful at solving the mental animation problems, and in fact were as successful as high-spatial individuals. In this case, the external visualization was augmented by the problem solver as a result his or her internal visualization processes. The augmented external display, in turn, relieved the participants of the need to maintain an internal visualization of the machine in memory.

In another recent experiment, Hegarty, Kriz & Cate (in press) examined the roles of mental animation and viewing an external animation in understanding how a relatively complex machine (a flushing cistern) works. One group saw a static diagram of the machine, another group saw diagrams of the machine in different phases of its operation and had to generate an explanation of how the machine worked (mental animation), a third viewed an external animation of the machine, and a fourth group both mentally animated and viewed an external animation. The results indicated that there were positive main effects on comprehension of both mental animation (generating an explanation from a series of static views) and viewing an external animation, and no statistical interaction between these two variables. This pattern of results suggests that the effects of internal visualization (in this case mental animation) and external visualization (viewing an external visualization) were complementary. As a result, those who first internally visualized and then viewed the external animation had better understanding than those who just mentally visualized or just viewed the external visualization. One possible interpretation of these results is the mental animation process induced individuals to articulate their intuitions about how the machine works, so that when they later viewed the external visualization, they could compare these intuitions to the actual physical process shown in the animation and pay particular attention to the parts of the process that they could not mentally animate.

In both of these examples, we see that there is an interplay between external and internal representations and processes. In the first example, the students augmented the external display to show the results of their internal visualizations. In the second case the external animation also augmented students' internal visualizations, such that they learned additional information from this animation that was not revealed by mental animation. A similar process of interaction between external and internal displays has been found when experts interact with data displays (Trafton et al. in press, Trickett & Trafton, 2002). These studies suggest that effective external visualizations are closely tied to internal visualizations and suggest that it is important to take account of internal visualizations in designing external visualizations.

Implications

In summary, the studies reviewed above suggest that rather than replacing the need to visualize internally, use of external visualizations is often accompanied by internal visualization activities, and effective use of these tools may even depend on a certain level of spatial visualization abilities. This suggests that educational studies need to be concerned with the development of student's internal visualization abilities, as well as being concerned with the development of the most effective visualizations.

Of course, one potential method of developing internal visualization abilities might be to expose students to external visualizations. For example, one approach might be to show external visualizations that model the internal processes that students are required to imagine in developing an internal visualization (Olson & Bialystock, 1983; Cohen et al, 2003). However, we need research on whether people can learn to internally visualize from external visualizations, and the conditions under which this learning takes place. Specifically, we need to know what types of interactions people need to have with an external visualization in order to learn from it. Given that some individuals may not appreciate how a visualization can be helpful to them (Cohen et al., 2003) we also need to understand what types of instruction to give students on how to interpret and use external visualizations. Importantly, the design of visualizations, and of instruction about how to use visualizations, should also take account of students internal visualization skills, as the research reviewed above suggests that the optimal form of instruction will be conditional on these.

More generally we need to understand more precisely how external visualizations augment internal visualizations. I have suggested one possibility, which is that they relieve the problem solving of the need to maintain complex internal visualizations in working memory. This is similar to the idea of distributed representations proposed by Zhang and Norman (1994). However, there are many other ways in which external visualizations might augment internal visualizations. For example, it is likely that there are limits on the complexity of the spatial transformations that can be imagined within the limited capacity of spatial working memory. External visualizations might augment these internal visualizations by presenting somewhat more complex transformations externally. However, this raises the question of how complex an external visualization people can comprehend, and the research presented above suggests that this might be somewhat related to internal visualization skills.

Finally in addition to studying how external visualizations augment internal visualizations, we also need to study how internal visualizations augment external visualizations. When a person mentally animates a static image of a machine (Hegarty 1992) or when a scientist or meteorologist mentally manipulates an external display of some data (Trafton et al, in press; Trickett & Trafton, 2002) he or she is augmenting the external static display with inner spatial transformation processes.

In summary, a review of the literature on understanding, learning from and producing external visualizations suggests that there is a complex interplay between internal and external visualizations, and that the design of effective external visualizations will be based on an understanding of internal visualization abilities. Rather than replacing the need for internal visualization processes, the development of technology to produce powerful external visualizations challenges us to better understand the nature of internal visualization processes, how to foster their development, and how they interact with external visualizations.

Acknowledgements

This research was supported in part by grants EIA-0313237 from the National Science Foundation and N00014-03-1-0119 from the Office of Naval Research.

References

Baddeley, A. D. (1986). *Working Memory.* New York: Oxford University Press.

Card, S. K., Mackinlay, J. D. & Schneiderman, B. (2001) *Readings in information visualization: Using vision to think.* San Diego, CA: Academic Press.

Bogacz, S. & Trafton, J. G. (2002). Understanding static and dynamic visualizations. In Hegarty, M., Meyer, B. & Narayanan, N. H. (Eds.) *Diagrammatic Representation and Inference, Proceedings of the second international conference on Diagrams, Berlin: Springer.*

Carpenter, P. A., & Shah, P. (1998). A model of the perceptual and conceptual processes in graph comprehension. *Journal of Experimental Psychology: Applied,* 4, 75-100.

Carroll, J. (1993). *Human cognitive abilities: a survey of factor-analytical studies*. New York: Cambridge University Press.

Chase, W. G. & Simon, H. A. (1993) Perception in chess. *Cognitive Psychology,* 4, 55-81.

Cheng, P. C-H. (2002). Electrifying diagrams for learning: principles for effective representational systems. *Cognitive Science, 26*, 685-736.

Cohen, C. A., Hegarty, M., Keehner, M., & Montello, D. R. (2993). Spatial abilities in the representation of cross sections. *Proceedings of the 25th annual meeting of the Cognitive Science Society.* Mahwah, NJ: Erlbaum.

Ekstrom, R. B., French, J. W., Harman, H. H., & Derman, D. (1976). *Kit of Factor Referenced Cognitive Tests.* Princeton, NJ: Educational Testing Service.

Ferguson, E. S. (1977). The mind's eye. Non-verbal thought in technology. *Science, 197,* 827-836.

Finke, R. A. (1989). *Principles of Mental Imagery*. Cambridge, MA: MIT Press.

Garg, A. X., Norman, G. R., Spero, L. & Maheshwari, P. (1999). Do virtual computer models hinder anatomy learning. *Academic Medicine, 74,* 87-89.

Gordin, D. N. & Pea, R. D. (1995) Prospects for scientific visualization as an educational technology. *The Journal of the Learning Sciences, 4,* 249-279.

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., Kriz, S. & Cate, C. (in press). The Roles of Mental Animations and External Animations in Understanding Mechanical Systems. *Cognition & Instruction.*

Hegarty, M. & Sims, V. K. (1994). Individual differences in mental animation during mechanical reasoning. *Memory & Cognition*, *22*, 411-430.

Hegarty, M. & Steinhoff, K. (1997). Use of diagrams as external memory in a mechanical reasoning task. *Learning and Individual Differences*, *9*, 19-42.

Hegarty, M. & Waller, D. (in press). Individual differences in spatial abilities. In P. Shah & A. Miyake (Eds.). *Handbook of Visuospatial Thinking*. Cambridge University Press.

Isaak, M. I. & Just, M. A. (1995). Constraints on the processing of rolling motion: The curtate cycloid illusion. *Journal of Experimental Psychology: Human Perception and Performance, 21,* 1391-1408.

Koedinger, K. R. & Anderson, J. R. (1990). Abstract planning and perceptual chunks: Elements of expertise in geometry. *Cognitive Science, 14,* 511-550.

Kosslyn, S. M. (1990). *Elements of graph design*. New York: W.H. Freeman.

Kosslyn, S. M. (1994). *Image and Brain*. Cambridge, MA: MIT Press.

Larkin, J. H., & Simon, H. A. (1987). Why a diagram is (sometimes) worth 10,000 words. *Cognitive Science, 11,* 65-99.

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

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

Mayer, R. E. & Sims, V. K. (1994). For whom is a picture worth a thousand words? Extensions of a dual-coding theory of multimedia learning. *Journal of Educational Psychology, 86,* 389-401.

Miller, A. I. (1986) *Imagery in scientific thought.* Cambridge, MA: MIT Press.

Narayanan, N. H., Suwa, M. & Motoda, H. (1994). A study of diagrammatic reasoning from verbal and gestural data. *Proceedings of the 16th Annual Conference of the Cognitive Science Society,* Lawrence Erlbaum Associates, pp. 652-657.

Nerssesian, N. J. (1995) Should physicists preach what they practice? *Science & Education, 4*, 203-226.

Nielson, G. M., Shriver, B. & Rosenblum, J. (Eds.) (1990). *Visualization in scientific computing*. Los Alamitos, CA.

Norman, D. A. (1993). *Things that make us smart : defending human attributes in the age of the machine*. Reading, Mass.:Addison-Wesley

Olson & Bialystock (1983). *Spatial Cognition*. Hillsdale, NJ. Erlbaum.

Pinker, S. (1990). A theory of graph comprehension. In R. Freedle (Ed.) *Artificial intelligence and the future of testing*. Hillsdale, NJ: Lawrence Erlbaum Associates, pp. 73-126.

Runeson, S. (1974) Constant velocity: Not perceived as such. *Psychological Research, 37,* 3-23.

Trafton, J. G., Tricket, S. B. & Mintz, R. E. (in press). Connecting internal and external representations: Spatial transformations of scientific visualizations. *Foundations of Science.*

Trickett, S. B., & Trafton, J. G. (2002). The instantiation and use of conceptual simulations in evaluating hypotheses: Movies in the mind in scientific reasoning. *Proceedings of the 24th annual conference of the Cognitive Science Society*. Mahwah, NJ: Erlbaum.

Tversky, B., Morrison, J. B. & Betrancourt, M. (2002). Animation: Can it facilitate? International *Journal of Human Computer Studies, 57*, 247-262.

Scaife, M. & Rogers, Y. (1996). External cognition: How do graphical representations work? *International journal of Human-Computer Studies, 45,* 185-213.

Schwartz, D. L. & Black, J. B. (1996). Shuttling between depictive models and abstract rules: Induction and fallback. *Cognitive Science, 20,* 457-497.
Schwartz, D. L. & Hegarty, M. (1996). Coordinating multiple representations for reasoning about mechanical devices. *Proceedings of the AAAI Spring Symposium on Cognitive and Computational Models of Spatial Representation*. Menlo Park, CA: AAAI Press.
Shah, P. & Carpenter, P. A. (1995). Conceptual limitations in comprehending line graphs. *Journal of Experimental Psychology: General, 124,* 43-61.
Shah, P., Mayer, R. E. & Hegarty, M. (1999). Graphs as aids to knowledge construction. *Journal of Educational Psychology*, *91*, 690-702.
Shepard, R. N. (1978). Externalization of mental images and the act of creation. In B. S. Randawa & W. E. Coffman (Eds.)*, Visual Learning, Thinking and Communication*. Academic Press, San Diego.
Shepard, R. N. & Metzler, J. (1971), Mental rotation of three-dimensional objects. *Science, 171,* 701-703.
Sims, V. K. & Hegarty, M. (1997). Mental animation in the visual-spatial sketchpad: Evidence from dual-task studies. *Memory & Cognition*, *25*, 321-332.
Spence, R. (2001). *Information Visualization*, Essex: ACM Press.
Zhang, J. (1996). A representational analysis of relational information displays. *International Journal of Human-Computer Studies. 45, 59-74.*
Zhang, J. & Norman, D. A. (1994). Representations in distributed cognitive tasks. *Cognitive Science,* 18, 87-122.