

External Cognition: how do Graphical Representations work?

Edited by Prof. Yvonne Rogers
and Prof. Anton B. H.
K

mikesc@cogs.susx.ac.uk

yvonner@cogs.susx.ac.uk

Abstract External Cognition: how do Graphical Representations work?

Keywords graphical representations, cognition, external and internal representations, interactivity, diagrams, animation, virtual reality, design guidelines.

E t r n a C o n t o n t o w o G r a p a p r s n t a t o n s w o r

an Yvonn o rs
n v r s t o f u s s B r t o n B H K

The Speaker, Betty Boothroyd, rebuked an M.P. for using a cardboard diagram in the Commons to explain overseas aid figures. She said "I have always believed that all Members of this House should be sufficiently articulate to express what they want to say without diagrams." [Guardian 7/12/94]

u ar

Advances in graphical technology have now made it possible for us to interact with information in innovative ways, most notably by exploring multimedia environments and by manipulating 3-D virtual worlds. Many benefits have been claimed for this new kind of interactivity, a general assumption being that learning and cognitive processing are facilitated. We point out, however, that little is known about the cognitive value of graphical representations, be they good old-fashioned (e.g. diagrams)

1 Introduction

Virtual reality and visualization¹, as means of representing and interacting with information, are very much at the forefront of technological development. An overriding intuition is that much can be gained computationally from interacting with virtual reality simulations or visualizing from three-dimensional dynamic images (Fairchild et al., 1993). Many benefits for industrial and educational applications have been claimed, such as powerful visualization tools for designers, architects and chemists (e.g. Earnshaw and Watson, 1993; Rheingold, 1991). However this is merely the latest in a long line of assumptions about graphical technological advancements, each claiming better ways of facilitating cognitive tasks. These include the ideas that:

- static pictures and diagrams are better than sentential representations²
- 3-D representations are better than 2-D ones
- solid modelling is better than wire-frame modelling
- colour is better than black and white images
- animated diagrams are more effective than static images
- interactive graphics are better than non-interactive graphics
- Virtual reality is better than animation.

Such generalisations about the benefits of advanced graphical technologies over good old fashioned representations, however, beg the question of what is actually gained cognitively from having more explicit, dynamic and interactive representations of information. Why, for example, should an animated diagram that changes in response to user interaction be more effective at facilitating problem-solving than static diagrams? Why not the other way round, where static diagrams are more effective than animations or non-interactive graphics

¹The term visualization is defined as “mechanisms by which humans perceive, interpret, use and communicate visual information” (McCormick and DeFanti, 1987).

² Epitomised by the widely-used proverb, ‘a picture is worth a thousand words’.

are better than interactive graphics and so on? Given this uncertainty, how can researchers and designers decide whether to take on board the immense cost and effort to develop a virtual reality application, for example, when a static diagram might be more effective for the task in hand?

The value of different graphical representations³, be they good old-fashioned or technologically-advanced, cannot be assessed adequately from our intuitions. To be effective a number of interdependent factors need to be considered, such as the level of experience with the graphical representation, the knowledge domain and the type of task. Whilst there have been numerous empirical studies investigating different aspects of graphical representations there has been little attempt to integrate the findings into an analytic framework. What is needed, therefore, is a more systematic approach for evaluating the merits of different kinds of graphical representations, one that is theoretically-driven and which accounts for the cognitive processing when people interact with them. Without such an approach we have no principled way of either making sense of the vast empirical literature on the benefits of graphical representations or of making predictions about the value of new forms, such as animation and virtual reality.

The current state of understanding is not encouraging. Most of the theoretically-based research on the role of external representations in cognitive science has been concerned with how we learn to read, write and understand written text. For graphical representations there is an obvious imbalance in terms of the work that has been done. On the one hand their value in helping to understand tasks/concepts presented verbally is well-documented but on the other there is no evidence of a detailed theory that explains this, e.g. *t r tur s*
ov r own wt wor nv st tn t ttv ts o p tur s on t t

³ Graphical representations include diagrams, maps, plans, animations and virtual reality and are distinct from propositional/sentential representations and formal notation(cf. Larkin and Simon, 1987).

operations on Annotated sources for on-line processing (Glenberg and Langston 1992, p. 129). In fact the quote applies equally to any learning or problem-solving situation utilising graphical representations whether text-comprehension is significantly involved or not.

Part of the problem may stem from the large variation in graphical representational forms, associated with a correspondingly wide range of functions. Past research spans a wide area from map design to technical illustration to the value of pictures for children learning science, with a *mélange* of methodologies, explanatory frameworks and mechanisms. Recent reviews are consistent in pointing out the lack of integration in the field. The problems here are severe for any attempt to provide an overall picture. For example, as Molitor et al. (1989) point out, a large number of studies have been concerned with the manipulation of task variables within highly-specific situations, reporting mainly on the success or failure of graphical representations to affect performance. As Winn (1993) notes, it is even difficult to make (practical) generalisations within this 'kind' of study, precisely because of their idiosyncrasies. In addition different authors have frequently ploughed their own furrow and have been highly selective in assimilating what other researchers have done. What has largely been absent, therefore, has been any attempt to explain how these experimental effects are produced psychologically, frequently ignoring recent work in cognitive science. Molitor et al. (1989, p. 27) comment that much of the empirical work on graphical representation has been *usuarily oriented to questions which not round on on-line processing*.

Why is there a lack of a suitable, explicit processing model? One reason may be that the form of graphical representation does not lend itself to systematic computational analyses. The theoretical frameworks and formal notations that have been developed for analysing verbal language are not applicable to the syntactically- and semantically-dense properties of graphical representations (Goodman, 1968). Another reason, as we argue later, may be that

there seems to be a pervasive (and possibly unwarranted) assumption that graphical representations must work in a certain way because of their figural nature. Thus many studies are almost 'black-box' in their approach to psychological mechanisms. Some, however, have attempted to look systematically at the effective perceptual features of graphical representations. For example Winn (1993) analysed diagrams in terms of a model of visual search, focusing on strategies for extracting information. His analysis identifies the importance of external features such as the spatial distribution and discriminability of elements of the diagram. He also points to important cognitive processes such as knowledge of content and symbol conventions in the reading process. Winn states that such accounts must as yet be insufficient. He points to a lack of graphical representation-specific research on search strategies but we would emphasise equally the paucity of work on determining how graphical representations are themselves represented and how this interacts with the kinds of high-level cognitive processes, such as applying knowledge of content, that Winn rightly emphasises.

We argue that an alternative approach is needed to understanding graphical representations: we need to ask what is the nature of the relationship between graphical representations and internal representations and to consider how graphical representations are used when learning, solving problems and making inferences. Such an enterprise means working towards a detailed description of cognitive mechanisms. In this respect we would point to must as ybe

with graphical representations and, hence, a useful account of their value. These models, thus do not match a second desideratum: an account which analyses more fully the interplay between internal and external when carrying out a cognitive task.

1.1 Internal and external representations

Within cognitive science, in general, there has been a move towards promoting the need to analyse the interaction between internal and external representations. In a special issue on situated action in the journal of Cognitive Science, Vera and Simon (1993) stress that, *A unitary perspective on external representations* (p12). Norman (1988, 1993) has for several years been describing cognition in terms of 'knowledge in the head' and 'knowledge in the world'. Larkin (1989) has also shifted her thinking from Larkin's and Simon's (1987) earlier computational model of diagram use - that focused primarily on internal representations - to considering the role played by external displays in cognitive problem-solving. Others, like Cox and Brna (1995) have been examining specifically the cognitive effects of external representations in reasoning tasks. External representations, here, may refer to both linguistic and graphical forms.

What we see emerging from this trend - to broaden and situate the base from which to explain cognitive behaviour - is external representations being given a more central functional role in relation to internal cognitive mechanisms. This is a substantial step away from traditional cognitive modelling and, significantly, an important theoretical advancement, that potentially allows us to account more adequately for how graphical representations work. Thus instead of trying to adapt internally-based processing models of cognition we can begin to specify characteristics of the internal/external relationship in the cognitive processing of graphical representations. The value of this is to focus our attention more on the *process* involved when interacting with graphical representations, the *properties* and the *functions* of different

graphical representations. In addition to enabling us to develop more appropriate cognitive models, we believe that this new perspective – which we have coined *transformation* – allows us to begin to assess more effectively how technological innovation in graphical representations should be approached.

In our examination of the emerging literature on internal/external representations we have abstracted three central characteristics which we consider as an useful analytic framework from which to explicate aspects of external cognition. These are *output economy*, *representational* and *representational*.

output economy - This refers to the extent to which different external representations reduce the amount of cognitive effort required to solve informationally equivalent problems. For example, Larkin and Simon (1987) point to the greater efficiency in geometry problem-solving for diagrams over sentential forms through their ability to provide direct perceptual recognition of geometric relations. Explicitly representing the problem state in diagrams in this way enables solutions to be more readily ‘read-off’. In contrast, solutions for the same problems represented as sentential descriptions typically are implicit and so have to be mentally formulated. This requires a greater computational effort.

representational This refers to how different external representations, that have the same abstract structure, make problem-solving easier or more difficult. For example, Zhang and Norman (1994) describe carrying out the same multiplication task using roman or arabic numerals. Both represent the same formal structure, but the former is much harder for people, used to working with the decimal system, to manipulate to reach the solution (e.g. LXV111 x X is much more difficult to solve than 68 x 10).

representational - this refers to the way graphical elements in a graphical representation are able to constrain the kinds of inferences that can be made about the underlying represented world. This characterisation is a term developed in recent work on the value of diagrams for solving formal logic problems by Stenning and colleagues (e.g. Stenning and Tobin, 1995; Stenning and Oberlander, 1995). A central idea is that the

relations between graphical elements in a graphical representation are able to map onto the relations between the features of a problem space in such a way that they restrict (or enforce) the kinds of interpretations that can be made. The closer the coupling between the

outline the relevant theoretical questions that need to be considered in understanding graphical representation applications and to suggest how they might apply to the design of innovative graphical technologies.

Before we move on to our analysis, however, we need to address potential sources of misunderstanding by considering the referential scope of key terms in our discussion.

A note on representation: static diagrams and animation in virtual reality

The term ‘representation’ has a variety of different meanings, depending on the context. A common distinction is between representation as process, and representation as product, as the outcome of this process. Process concerns the transformations and preservations that occur in deriving the representation from what is being represented. Description of product is typically concerned with structural characterisations of the representation, for example as image-like, mental model or propositional. Confusion might arise since the two senses, process and product, may be used interchangeably. In fact the two cannot always be easily separated, since characterisation of structural properties is usually related to a particular processing model. Here we shall discuss representation in both senses.

The classes of static diagrams and animations are considered distinctive, in so far as they have been identified as having different characteristics in the literature. It is acknowledged, however, that there is likely to be some overlap between what constitutes a diagram and an animation, especially for displays that comprise of both static and animated components. There are many different exemplars of diagrams and no single accepted taxonomy that can be conveniently employed to describe them, although there is good evidence emerging for stable classification strategies (Cox and Brna, 1993; Lohse et al., 1991). There may, in fact, not be a single, criterial feature for the term ‘diagram’. Some authors seek to draw a

distinction between representations like graphs, describing quantitative data in two dimensions, and other, less-constrained types. We would probably subscribe to that view but it is not crucial for us here. We would rather adopt a position similar to that of Winn (1987, p. 153) who treats diagrams as representations with the function of being *to* *s* *w* *o* *p* *r* *s* *s* *n* *s* *t* *r* *u* *t* *r* *s* *o* *t* *n* *t* *v* *s* *o* *r* *t* *o* *p* *t* .

Animations are equally difficult to define and, again, there is, as yet, no single theoretically- or even empirically-grounded classification scheme available. Animations - be they computer, film, video or other media-based - differ from static diagrams in presenting a series of rapidly changing static displays, giving the illusion of temporal and spatial movement. This can be achieved through a range of techniques. For example in 'multi-dimensional' animation interdependent objects appear to move in relation to each other; in 'partial' animation certain parts of a display move whilst the rest of the display remains static; in 'artificial' animation implicit movement is made explicit or processes normally invisible to the eye are made visible. While not an exhaustive classification, we can see the diversity of animation 'types'.

The third class of graphical representations that we examine is virtual reality or virtual environments. These are computer-generated graphical simulations, intended to create *t* *u* *s* *o* *n* *o* *p* *r* *t* *p* *t* *o* *n* *s* *n* *t* *t* *n* *v* *r* *o* *n* *n* *t* *r* *t* *r* *t* *n* *t* *r* *n* *s* *r* *v* *t* *o* *n* *s* *u* *n* *n* *v* *r* *o* *n* *n* *t* (Gigante, p.3, 1993). Images are displayed stereoscopically to the user, via a head-mounted display. Objects within this field of vision can be interacted with via a dataglove or other input device, use of a virtual reality headset can change the field of vision in the virtual world and users can 'fly' around the virtual world through gesturing. A major motivation for virtual reality systems is to enable people to become 'immersed in the experience' of interacting with external representations (Kalawsky, 1993). However this is difficult to operationalise (Sheridan, 1992) and there is no taxonomy of types of virtual reality immersion. Most virtual reality classifications are based on the types of graphical

techniques used for rendering 3-D objects and in terms of applications that may benefit from being represented in virtual reality (Kalawsky, 1993).

E p r a w o r d s o n r a p a r p r s n t a t o n s n v o v n a r a s a n a n a t o n s

Having set out some desiderata for a study of graphical representations we will now try to make our ideas more concrete. Rather than attempt a global review, we shall concentrate on a small number of influential studies to draw out some general issues pertinent to our aim of assessing the pros and cons of different kinds of display. We shall examine two studies that concern the use of static diagrams and two that have investigated animated displays. These have been chosen as examples which have clear aims to show how graphical representations might work and what processes are involved. We shall adopt the format of first describing the findings and then offering a critique before making some general comments on the theoretical issues that surround graphical representation research.

1 s a r a o n s t a t a r a s

Work on static diagrams represents a considerable corpus of research from which it is hard to make generalisations. Winn (1987), reviewing the field, notes that there is an interaction between (at least) ability level, diagram format and task type to be considered in drawing conclusions across studies. We shall consider here two studies that have looked at the value of diagrams for problem-solving: Larkin's and Simon's (1987) study of physics and geometry problems and Bauer's and Johnson-Laird's (1993) study of logic problems.

Larkin and Simon (1987) analysed examples taken from classic physics (pulleys and weights) and geometry (theorem proving) textbooks. Their aim was to develop computational models that allowed a contrast between processing of 'sentential' and 'diagrammatic' representations which contained the same information about the problem. In

the first case elements appear in a single sequence, while in the second they are indexed by their location in two-space. Their theoretical analysis suggests that a diagram *pr s rv s* *p t t n or t on* *out t topo o n o tr r t ons on t* *o pon nts o t pr* *w t s nt nt r pr s nt t on o s not* (p.66).

The approach taken by Larkin and Simon provides an explicit formalism. The elements of their system are (i) data structures that represent the problem to be solved (ii) productions that contain knowledge of the laws of the domain (the 'program') and (iii) an attention manager. They propose that a diagrammatic data structure may differ markedly from an informationally-equivalent sentential one through affording the possibilities of easier search.

They observed that when data structures are informationally equivalent *In r n s r n p n nt o r pr s nt t on* (p. 71). Bauer and Johnson-Laird (1993), however, postulated that for certain kinds of problems, diagrams should help reasoning, a claim based on Johnson-Laird's (1983) model theory of deductive reasoning. They investigated the role of external representations, in the form of schematic diagrams, on the solving of deductive reasoning tasks. The problems were double-disjunctive reasoning, which require reasoners to keep track of various alternative states in order to solve them. Because of the difficulty of taking into account many models of the premises subjects are known to perform poorly on these types of problems. Bauer and Johnson-Laird hypothesised that providing diagrams should enable reasoners to keep track of alternative models.

Figures 2, 3 about here.

Bauer's and Johnson-Laird's (1993) first reported attempt at developing a schematic diagram to make explicit the alternative possibilities was largely unsuccessful. This they attributed to their using arbitrary and abstract icons for representing explicitly the alternatives, which were found to be of no help to the reasoner (see Figure 2). Their second attempt, however, was more successful. Two types of more concrete diagrams were constructed: one based on an electrical circuit and the other a jigsaw. In both examples, a particular problem-solving context was provided from which to make the deductions. The instructions for the circuit representation of the problem was couched in terms of switches and lights being on or off in the circuit whilst the instructions for the jigsaw representation were expressed in terms of completing a path from one side of the figure to the other. This involved inserting shapes, corresponding to specific people specified in the reasoning problem, into slots in the path, corresponding to particular places (see Figures 3a & 3b). In both examples, therefore, the subjects were required to solve the reasoning task by mentally transforming parts of the diagram. In doing so, the solvers no longer need to solve the problems entirely in their head but can work them out by interacting with the diagrams.

Indeed, the results showed that performance was significantly better and faster when using the diagrammatic representations than when solving the same problems using sentential representations. The findings seem to provide further support for the important role of diagrams as external memories, enabling a picture of the whole problem to be maintained simultaneously, whilst allowing the solver to work through the interconnected parts (cf. Larkin, 1989; Larkin and Simon, 1987; Zhang and Norman, 1994). Although it could be argued that a sentential representation can also act as an external memory aid, the extent of the ‘computational offloading’ is considerably less. The reason being that the problem states and its solution are more explicitly represented in the diagram than in the sentential representation, meaning that less inferencing is required.

Both of the above studies show the potential value of graphical representation for aiding problem solution in terms of search, recognition and inference. However we need to ask the related questions: how much light do they shed on the role of the external representation and how does this mesh with details of internal representations (cognitive mechanisms)? Consider first the Larkin and Simon (1987) account. Their principal concern is with diagrammatic internal representations, which as Parkes (1993, p. 37) points out provide

... *... to the proportion of the total work done by the representation of the problem state.* Larkin and Simon (p. 66) describe such representations as having the property of corresponding

... *... to the problem state.* However, their account leaves open the question of (i) how this is produced in human beings and (ii) what work is being done by the external and the internal representations respectively. Consider the following quote (p. 92):

... *... to the problem state.* But we ... *... to the problem state.* It is hard to get a precise understanding of how ‘perceptual inferences’

differs from the model theory by postulating that representations support the construction of a mental model (p.2). This is achieved through the information available in the diagram restricting the possible interpretations of the problem and in so doing guiding the reasoner to make the correct solution. Thus, certain diagrams are more effective than others because they exploit better the constraining properties of varying graphical forms. For example, Stenning and Tobin (1994) claim that Euler's Circles (see Figure 4) are more effective than 3-D cube diagrams in helping subjects solve logic problems because the geometrical constraints of the intersecting circles represent the logical constraints much better. In other words, a diagram is more likely to afford a particular reading of the problem and way of solving it than a sentential representation because it is less expressive (i.e. decreases indeterminacy). Having built a mental model of the combined external representations (the instructions and the diagram) that satisfies the premise in their own minds, it is unlikely that the solvers will then build alternative, but equally plausible, models of the problem (Cox and Brna, 1995).

Figure 4 about here

Conclusions

One of the strengths of studies, such as Larkin and Simon (1987), is the postulation of an explicit model of cognitive processing. However, in reviewing research into the role of animations in learning and problem-solving contexts we failed to find any similarly detailed models. Thus, in our attempt to consider in more detail how animation is processed we decided to critique two empirical studies that sought to investigate the mechanisms by which animations are effective in making inferences from graphical representations of physical systems. These are Hegarty (1992), which focuses on mental animation and Kaiser et al. (1992), who focus on external animation.

In Hegarty's study the graphical representations used were static canonical diagrams of

pulley systems whilst in Kaiser et al.'s study, both static and animated canonical graphical representations were used to depict objects falling, being severed, or being displaced from various dynamical systems (e.g. pendulums and moving planes). The primary aim of Hegarty's study was to ascertain the extent and form of mental animation that occurs when making judgements about the motion of pulley systems. In contrast, the main objective of Kaiser et al.'s study was to determine how external animations enable more effective judgements to be made about the trajectories of moving objects compared with static diagrams. In both studies, subjects' were required to reason through initially comprehending a verbal problem together with a static or dynamic graphical representation used to convey the problem state, and then predict correctly future states or trajectories of part of the system depicted in the graphical representation.

Hegarty's central idea is mental animation, which involves simulating mentally, in a serial manner, components in the graphical representation of the pulley system. An obvious reason for this is that we are unable to animate all parts of the diagram at once, due to the constraints of working memory. It also seems plausible, given that we can only perceive the working of certain aspects of a real world pulley system at any one time, depending what is in our field of view at that time. With real-world pulleys, however, the motion of each part is always available; we need only to follow the way the components move to make judgments about them. Moreover, we can do this in a haphazard way. With diagrams, however, Hegarty argues we make inferences about the motion of the static parts by following the temporal order of the causal chain of events from input to output.

This level of theorising seems intuitively plausible for explaining how people reason with relatively simple pulley system problems and is to some extent supported by her empirical findings. For more complex systems, Hegarty suggests that other mental strategies are likely to be used. However, the form that these alternative forms might take, how they develop and whether they are used in combination with mental animation or separately is

beyond the scope of her paper. Likewise, the actual functional role of the graphical representation is not discussed in her theory of incremental animation, although she does acknowledge that it needs to be researched further.

Figure 5 about here.

In contrast to Hegarty's approach, Kaiser et al. (1992), explain reasoning about mechanical systems in terms of what the external representation does for the learner. Like Hegarty, they stress the importance of information being processed sequentially, but in terms of the external representation being able to 'temporally parse a multi-dimensional problem into unidimensional components' (p. 671). In doing so, they propose that the distinct state changes that have to be recognised to make correct judgements about the system are made more obvious through an animation than with a static display. The idea that the external representation does the 'temporal parsing', rather than the problem-solver having to do it, is illustrated with an example of common-sense reasoning about the C-shaped tube physics problem (based on McCloskey et al., 1980). The main finding is that when the problem is represented as a static 2-D representation (see Figure 5), students often incorrectly infer that the projected motion of the ball on exiting the curved tube continues in a curvilinear trajectory. Kaiser et al. found the same effect for both free choice and forced choice conditions. However, when shown various incorrect and the correct animation sequences in a forced choice situation, students invariably selected the correct 'straight' trajectory. Kaiser et al. explain this performance shift in terms of the animation temporally segregating the ball's behaviour from the rest of the scene.

diagram research. In particular they assume that running mental models and parsing external animations use the same structures and functional processes as when perceiving real-world dynamic systems. For example, in an earlier study Hegarty et al. (1988) proposed that people decided which attributes of a system were relevant to judging mechanical advantage on the basis of 'causal models' of mechanical systems arising from relevant (physical) experience with such systems. The appeal to an equivalence in processing, however, does not help us in understanding the merits of different forms of graphical representation in terms of how they are processed and interacted with for various tasks. A more pertinent question to address, then, is how do we understand and make connections between the static and animated forms that represent the dynamic processes of real systems such that we can make inferences about them both?

The same lack of specificity is seen in Kaiser et al.'s (1992) explanation of the superiority of animated over static forms in terms of the visible temporal parsing of the ball in the container system, making change in states more obvious. An alternative explanation for the difference in performance could be in terms of experience with the two representational formats, reflecting more a difficulty with interpreting the canonical forms in the diagram in relation to the problem that had to be solved rather than one of not being able to recognise the significance of the temporal parsing of the objects in different states. This objection reiterates our concern that there is a crucial role for expertise and practice that is not being recognised. Most importantly, whilst providing further support for the value of the explicitness inherent in animations, Kaiser et al.'s study offers no explanation of the cognitive mechanisms involved in learning and reasoning with animations. Pedagogically, too, it is unclear how animation can facilitate learning or problem-solving. In particular, we would emphasise the absence of an analysis of how a better level of understanding can result from seeing objects moving explicitly as opposed to having to imagine how they move. Indeed, Kaiser et al. comment on how subjects who had been shown the animation first and then a static diagram of the same problem performed no better than those who had

just been shown the static diagram. Here again, we have further evidence that the benefit of viewing an animation is transitory and not readily mapped onto the static representation with its more arbitrary conventions.

Conclusions for students

In all of these four studies, then, we are left with questions about the mechanisms by which diagrams and animations are effective. How do viewers identify the 'key' features and constraints of a graphical representation and then map them onto the relevant aspects of the problem to be solved? Further, there seems to be a real issue about what problem the subjects were 'really' solving and how far expertise and experience are central factors. It is often hard to separate general claims about graphical representations per se from factors that have to do with individual differences in ability in the subject or understanding of the domain-specific genre of the diagrams involved. In domains with highly evolved notations such as geometry or physics, diagrams are not 'merely' aids for solution but play an essential part in the process of knowledge acquisition and depiction. A circuit diagram, an architectural plan or a mathematical notation comprise a set of meaningless symbols to the uninitiated; they only take on their intended meaning through learning the conventions associated with them. In a real sense it is impossible to develop expertise in these subjects without the ability to both read and produce diagrams of a particular sort (cf. Anzai 1991). This strongly suggests that - in such domain-specific cases at least - diagrams can only trade on established domain knowledge to be effective. The point is well-put by Larkin and Simon (1987, p.71): *I students produce or not a solution or not without a representation.*

4. Process and answers

The studies we have reviewed are, of course, a tiny selection from a vast range. However they serve to illustrate two of our major themes, that there is a lack of an adequate cognitive processing model and that focusing on the externality of graphical representations to see

how they work is of crucial importance to a better understanding. Below we examine further why the former is so and contrast this with a more detailed analysis of the few studies that have begun to analyse external/internal representations in cognitive processing.

4.1 **Pro ss n an t r s an afa**

One problem with the Kaiser et al. (1992) study noted above was the lack of explanation of how subjects recognise the temporal segregation of the objects as being central to an understanding of mechanical systems. This seems to be because the external and internal representations are assumed simply to have the same characteristics. This is an example of what we shall call the 'resemblance fallacy', which has a much wider appearance in the graphical representation literature and may help to explain something of the apparent unwillingness to specify processing models. It is prevalent, we believe, because the structure of graphical representations, their spatial/iconic/figural qualities, promotes an *intuition* as to their value as an input for perception/cognition whereas the reality is that we have no well-articulated theory as to how such an advantage might work. Evidence for our over-reliance on such intuitions can be seen by examining the kinds of arguments that have been made for the links between graphical representations, perception and internal representations. The possibility of different representational formats - image/proposition/mental model - and their properties interacts with the issue of how graphical representation might work.

Consider first, work on conventional, static graphical representations. By design these are

example Winn (1987, p.159) summarises the relevance of work on imagery to graphical representation thus: “ *s stu s p o o r s r t t s t o t o o w n*
o n u s o n G r p o r s n o u r s t u n t s t o r t n t s t t n t u r n t
s r o r t t o r n r t n t p s o t r And Reed (1993, p.299) claims a
s t n t s r t t w n t u n t o n q u v n o p t u r s n s, stating
 that: “ *w o u v t t r u n r s t n n o o w s p r s o v n w*
t t r u n r s t n n o o w p t u r s p r s o v n

The problem with this line of argument is that it does seem to rest on intuition. What can ‘encourage’ and ‘easier’ mean in terms of mechanism? Further, while pictures can undoubtedly serve to stimulate imagery under certain circumstances (e.g. Finke, 1990) it is by no means clear that they are necessarily represented in this way. Halford (1993) points out that we do not have to accept any more than a mapping between relations for an external representation-internal representation pair. In addition there is some doubt about the extent to which imagery is computationally important and processing may be better explained in terms of other representational forms (e.g. Anderson, 1990; Molitor et al., 1989; Pylyshyn, 1973). In short the case for an intimate relationship between graphical representation and images may not be logically compelling and is currently heavily under-specified.

by no means clear that stimuli surprised in recent years. These studies extended in that they... runnily u rn, ma

Similarly Bauer and Johnson-Laird (1993) talk about the *input-output* relationship. Thus, while the mental model resulting from interaction with graphical representations need not be thought of as image-like (cf. Glenberg and Krulley, 1992), there is often an apparently close relationship posited between the two and, hence, the possibility of another intuitive and unsubstantiated link between picture input and representation type. Molitor et al. (1989, p.10) describe the situation thus: *Input-output relationships are not necessarily linear or direct. However the same authors also note that the mental model construct is virtually unnecessary for learning* (p. 10), a conclusion echoed by Hong and O' Neill (1992). As with images such a range of possibilities argues for caution in positing a necessary relationship between the pictorial nature of graphical representation and a particular representational format. At the very least, different graphical representations and/or tasks may engender different kinds of representations and this must remain an issue for future research.

The discussion so far of the resemblance fallacy has been exemplified by work done with static diagrams. However there is evidence that some of the same kinds of problem are occurring with work on animations. Here, as we noted previously, there is an assumption that 'adding' animation to an equivalent static display will be advantageous. But why should this be so? In articles about animation we commonly find an intuition-led chain of assumptions, echoing the same causal chain of reasoning used to account for the efficacy of diagrams (external representation producing a mental image or mental model which in turn results in better learning or reasoning) that was criticised earlier. Here, an illustrative line of reasoning goes something like this: animations can show motion explicitly and 'directly' and hence provide more accurate information (Kaiser et al., 1992); this reduces processing demands on working memory allowing other tasks to be performed (Rieber and Kini, 1991) and enables more 'useful' mental models to be formed for solving problems (Park and Gittleman, 1992); these in turn facilitate learning or reasoning. Whilst the first part may be factually accurate, the rest does not logically follow.

As argued above in relation to diagrams, we cannot simply assume a privileged relationship between a graphical representation of a system - in this case an animation - and someone's understanding or ability to reason about it, by virtue of its resemblance, albeit highly simplified and schematised, to the dynamic properties of a real-world system. As with diagrams used in specialised domains, e.g. physics or geometry, a person has to learn to 'read' and comprehend the significance of the content of the animations in relation to other information that is being presented verbally or as text and to assimilate this to their current understanding of the domain. This requires making multiple connections between what the animations are intending to convey and the abstract concepts that are being learned about. How students integrate information arising from different representations of knowledge is crucial (Laurillard, 1993).

4 Pro ss n an t nt rna t rna

The force of our comments, however, is not solely to do with being less intuitive in our accounts. Consider the claim by Larkin and Simon (1987, p. 97) that:

“ nt r t us so r s n ot rp tor r pr s nt t ons n u n

as mentioned previously, to be more of a central concern in cognitive science. Larkin (1989) has tried to tackle the problem by outlining a computational model called DiBS, that represents information available in external displays as data structures that enable internal operators to be cued as to know what to do next. The model's central searching mechanism is based on the observation that 'each step requires only looking at the display, and doing what it suggests, without more effortful mental calculation or storage' (p319). DiBS, therefore, works largely by manipulating attributes of the external display. The examples that Larkin has chosen to represent in her model are well suited to the transformation of external data structures. They include simple everyday problems (e.g. brewing coffee) and textbook problems (e.g. linear equations) that once learned become highly routinised and error-free. Hence, for these kinds of tasks there is no need to activate any internal representations other than a very general mechanism that is characterised as knowing 'where an object wants to go' in each step of the task. As such, DiBS is a model of how external representations cue a set of task procedure-based actions for solving a task. The external mechanisms are largely

inspired by Larkin's (1989) 'move to the external', number of empirical studies have been 0.084 Tc 24 TL0.6770 Tw e

How the gaps are filled and what cognitive mechanisms are involved, however, is not clear. This is also true of Zhang and Norman's (1994) recent analysis of distributed cognitive tasks (between internal and external representations), where they argue that an analysis of the relation between different external forms of representing the same abstract problem (in this case the Tower of Hanoi) is necessary before considering the processes that are activated when solving the problem. The study was designed so that in certain conditions subjects had to internalise several rules to carry out the task whilst in others the same rules were embedded in the external display. Their findings indicated that the fewer rules subjects had to internalise the easier it was for them to perform the problem-solving task. The implication is that external representations can significantly change the nature of a task through constraining the permissible moves allowed in solving the task. Furthermore, this form of 'computational offloading', i.e. implicitly embedding rules in the external representation as opposed to making subjects internalise them, is thought to reduce the load on internal working memory providing more 'space' for planning subsequent moves. Although Zhang and Norman did not investigate the processes involved in interacting with graphical representations in their experiment, they do suggest that the nature of the relationship is likely to be uni-directional: *processes represented by external representations* and that moreover *processes represented by external representations* (p118). We would argue, however, that the interplay between internal and external representations in problem-solving is likely to be more complex, involving cyclical, interacting processes, especially when considering how graphical representations are both perceived and acted upon.

Before moving onto our final discussion of how we can consider the design of graphical representations based on an analysis of the relationship between internal/external representations we briefly introduce our third category of external representations, virtual

may prove to be as incorrect. For example, results from a recent study investigating transfer of training in virtual reality systems found that subjects learnt performance characteristics specific only to the virtual reality context, which were of no use when carrying out the same task in the real world (Kozak et al., 1993). These preliminary findings suggest, therefore, that the actual experience of being immersed in a virtual reality world is quite distinct from interacting with real world artefacts. The value of virtual reality, therefore, should not be

what aspects should be omitted and what additional information needs to be represented that is not visible in the real world but would facilitate learning. From a cognitive perspective, it enables us to assess the benefits of virtual reality in terms of the processing mechanisms that operate at differing levels of abstraction of information. For example, we can analyse differences in task demands and performance characteristics for specific tasks, e.g. taking off or landing for different virtual reality simulations, ranging from presenting simple canonical structures (e.g. schematic outlines) to more fully rendered depictions of scenes. Hopefully, this way the pitfalls of the resemblance fallacy can be avoided.

Another way in which the notion of virtual reality immersion has been characterised is in terms of 'steering' the interaction. Here, the intuition is that virtual reality simulations provide more opportunities to visualise and manipulate the behaviour of abstract data structures or processes which are not normally visible to the naked eye. For example, NASA have developed a Virtual Wind Tunnel, whereby a scientist (who is a computational fluid dynamicist) controls the computation of virtual smoke streams by using the finger tips (Gigante, 1993). Abstract equations for the computed airflow around a digital model of an aircraft are translated into visible smoke streams. By moving around the virtual aircraft in the virtual reality environment and visualizing the smoke streams, the fluid specialist is

simulation should be analysed, therefore, in relation to how it integrates with ways of interacting with other existing forms of external representations in professional practice.

In/on representations

At the beginning of this paper we asked the important practical question of how designers could determine which kind of external representation to choose from - be it text, diagram, multi-media or virtual reality - for the domain or task they are designing for. As has become clear through the paper, however, to answer this question depends on having a better understanding of internal representation/external representation interactivity. In addition, it requires addressing specific issues, such as what form the display should take, what information should be made explicit, how this should be represented, how this maps onto the object/concept being represented and which graphical style to use. Our analysis of the graphical representation literature across a variety of disciplines, however, led us to the conclusion that despite a plethora of empirical studies on how different graphical representations affect performance and a few theoretical analyses, the findings are difficult to generalise beyond the specific features investigated in each study. Moreover, the majority of studies were largely silent about the criteria used in the design (sic) of the graphical materials for the experiments. It appears, therefore, that we need to be more explicit about the selection and design of graphical representations for both applications and experiments investigating cognitive aspects of interacting with graphical representations.

Attributes

As Bauer and Johnson-Laird (1993) discovered, facilitation of problem-solving depends on the kind of graphical representation being used. So what are the attributes of a 'good' graphical representation? The issue of 'good design' has been the subject of a number of different studies which have attempted to give guidelines for producing graphical displays

(e.g. Goettl et al. 1991; Kosslyn, 1989) but these have been concerned primarily with the

diagrams is clearly important for their effective use but it may also be that experience with static forms itself may be a useful precursor to the ability to read more dynamic ones.

Under such circumstances we need insight into how people read and interact with diagrams. The stress here is important for our ideas about developing good diagram skills. In the vast majority of studies and analyses of static diagrams the assumption has been that the subject does nothing to change the external form. This may well be true for cases such as library books or slides in a presentation but it may not be the optimal case. Koedinger and Anderson (1990) observe that high school students frequently use annotations to problem diagrams to hold together information needed for inferences. Such a strategy may also be performed mentally but, given previous arguments, is probably more efficient for learners when external. In fact the logical conclusion of these arguments is to maximise the load on the external representation. As has been observed in several cases making a 'cognitive trace' available for problem-solving is of great benefit (e.g. Merrill & Reiser, 1993). One lesson for diagram use might, therefore, be to promote opportunities for external manipulation (i.e. cognitive tracing) as well as encouraging production skills, as we noted above.

Good diagram design also has the crucial requirement that the degree of abstraction of material should be appropriate to the varying demands of the task and learner's ability. Levonen and Lesgold (1993) describe SHERLOCK, a computer-based electronics coaching system, which has the facility for representing both realistic (picture-like) diagrams of the system and schematised expert representations of the same domain. Switching between the two enables a kind of apprenticeship learning. Consistent with this approach Cheng (1993) advocates the availability of multiple representations, from specific examples to overviews, which learners could choose to look at as they wished. While this raises issues about integration between different views it also emphasises the importance of learner control. There is no reason to doubt, for many situations, that multiple representations could be made available within a 2D diagram. However it may also be that this is something that may

well be better achieved in other forms.

Conceptual issues and future viewpoints

The above discussion points to a number of factors which designers should be aware of. We suggest that it is useful to begin to formalise them as a set of general conceptual design issues, akin to the set of cognitive dimensions of notations that Green (e.g. 1989, 1990) has advocated, for describing important features of the design of programming languages and software tools to support users' tasks. Firstly, they can help bridge the gap between our conceptual understanding of how graphical representations work and the practical concerns of designing graphical representations. Secondly, requirements for future technological developments also can be assessed in relation to cognitive processing. Thirdly they can help us reframe design questions. We could ask what is required to design advanced graphical representations that can be of 'added' cognitive value for particular users, domains and tasks? Below we present an initial attempt to identify some of the key conceptual design issues.

Expressiveness

Diagrams, animations and virtual reality can in their respective ways all make salient certain aspects of a display. A design objective, therefore, should be to facilitate perceptual parsing and inferencing, through directing attention to key components that are useful or essential for different stages of a problem-solving or a learning task. In addition, the various graphical representations can represent 'hidden' processes which underlie complex phenomena. The aim, here should be to facilitate higher level understanding, i.e. cognitive inferencing but also in relation to how this interacts with perceptual processing. As with 'perceptual inferences', the users may need much prior knowledge in knowing how to interpret what is shown.

Contraints and interaction

Diagrams that have been already constructed allow the user to leave cognitive traces, i.e. mark, update and highlight information. However this is a limited function. There is no

possibility of interaction or feedback - the user cannot test new configurations. In contrast, when interacting with animations and virtual reality objects there is more scope for providing feedback but less for leaving cognitive traces. For example, various parameters of a computer-based model can be set in a virtual reality or 3-D simulation (cf. microworlds) and the outcome directly observed. Graphical representations should be designed with a view towards how they support different kinds of cognitive tracing and levels of interactivity.

Eas e of p r o d u c t i o n

Related to the above issue is ease of production of a graphical representation. It appears, that diagram production and comprehension are intimately related. A history of being taught to draw diagrams makes for fewer problems with understanding new ones. This is particularly important for domains where evolved notations are crucial. However, where the possibility of acquiring expertise is limited, the demands of reading the diagram efficiently may be too great. Recent software developments now make it possible for users to select alternative or partially animated views of the same process, and to play (and replay) them at different speeds, thus enabling multiple abstractions to be interpreted. Furthermore, software is being developed that will allow novice users easily to construct their own animations through compiling components from a toolbox of animations or modify pre-designed animations. The hypothesis about diagram production and comprehension could be tested for these more interactive forms of animations: having a better understanding of how to create animations will enable people to have a better understanding of how they work and what they are trying to convey.

C o n v e n t i o n s o f r e p r e s e n t a t i o n s

The conventions of constructing 2D diagrams have largely evolved to be complementary to textual expositions. In some cases it may be that text is indispensable for understanding the function of a particular diagram. However fairly mundane factors such as spatial separation of text and diagram may significantly increase the computational load involved in comprehension (e.g. Sweller et al., 1990). In contrast, animations and virtual environments

have been designed to be largely graphical, although they may be accompanied by spoken narration or verbal text. Studies have shown that it can be more difficult to integrate written text with these kinds of graphical representation than with static diagrams. For example, response times from Rieber's (1989) study of combining text with animations to represent Newton's Laws of motion indicated that the subjects simply viewed the animations and then moved immediately onto the next screen of information without reading any of the accompanying text. Other studies which have combined spoken narration with animations, however, have fared better, showing that this combination is more effective. For example, Mayer and Anderson's (1991) study of subjects' understanding of the operation of a bicycle tyre pump, showed comprehension to be better when the information was depicted as an animation with concurrent narration, than when presented just as an animation. Having parallel auditory narration could also be effective for virtual reality to guide users in exploring and interacting with the environment. Hybrid graphical representations could also be developed that allow users to interact with static diagrams on a computer display by adding animations or conversely, allowing users immersed in a dynamic virtual reality environment to interact with static objects (e.g. jotting notes into a virtual notebook). The objective, here, would be to provide support for different kinds of interactivity.

Distributed representations

While we have not commented on this aspect in the paper, diagrams offer the possibility of joint evolution of representations, e.g. in idea-sketching, where planning can be facilitated using notations in any framework that suits the task. Here temporary conventions can be set up, which also has the drawback of being potentially unintelligible at some later date or to others. Various shared drawing tools also have been developed to support collaborative sketching and designing (e.g. Ishii and Kobayashi, 1992; Scrivener et al., 1992). Collaborative design sessions can be recorded which when played back re-construct the collaborative drawings as animations. virtual reality environments can also provide opportunities for virtual construction of graphical representations for users in geographically dispersed locations. However, the value of enabling collaborative construction and editing

of graphical representations in terms of enhancing task performance is only beginning to be researched.

– **CONCLUSIONS**

A major aim of this paper was to examine the strength of claims for the value of advances in graphical technology for facilitating cognitive tasks. We have seen that these claims are often underpinned by assumptions which have little empirical support and/or insufficient theoretical grounding. In addition there has been little progress towards a framework, either methodological or theoretical, that might allow the designer to produce and evaluate new forms of graphical representation or even improve on existing ones.

Part of our argument has been to pick over the bones of previous studies to see why there has been so little progress and/or integration despite an enormous volume of research. The answer seems to be in several parts. Firstly, the studies have been highly detailed and do not generalise. Secondly, they have failed to produce adequate rationales for the material tested, making it difficult to determine what is actually being assessed. Thirdly they make assumptions about the kinds of linkages between external representation and internal representation which are rarely articulated (cf. the resemblance fallacy) or, if they are, may not give sufficient weight to the role of the external. Fourthly, articulating the links require theoretical analyses, of which those that might seem appropriate, are only beginning to emerge as theoretical developments in cognitive science (e.g. Hutchins, 1995).

Most existing accounts of how graphical representations are effective, therefore, have been black box in nature – there exists a gap in terms of explaining adequately any cognitive processes involved. For example the account of internal process may be couched in terms of ‘applying knowledge of content’ but give us little about what kind of internal representation is mediating task performance. Likewise, as we argued in describing the resemblance fallacy, making assumptions that the internal representation is a mental model or image-like

may simply give the illusion of solving the processing-internal representation-external representation riddle. But instead, the problem of explaining the value of graphical representations is shifted simply from an external to an internal account. In contrast we promote an alternative approach that analyses how different graphical representations work in terms of core 'external cognition' processes and properties of the graphical representation, e.g. computational offloading, re-representation and graphical constraining. We believe that such an enterprise is central to evolving a more adequate account of the cognitive benefits and mechanisms involved.

Related to this is a further, critical and under-acknowledged theme, that of interactivity. Specifying how people interact with graphical representations, when learning, solving problems and making inferences, is complex since it will involve not only a specification of the cognitive mechanisms alluded to above but also some sense of the behavioural aspects. For example the fact that students prefer to mark diagrams as they work, the established value of cognitive traces and the dialectic between graphical representation production and use all point to a need to conceptualise graphical representations as more than passively observed, with obvious implications for design and innovation. In turn the potential significance of such activity will be a function of variables such as the level of experience with the graphical representation and knowledge domain, type of task and abstractness of information being represented. Many of the presumed benefits of good-old fashioned graphical representations (i.e. static diagrams) were considered to be due to years of practice of perceptual processing of visual stimuli and the learning of graphical conventions. This may help us to understand why advanced graphical technologies (e.g. animations and virtual reality) have not, as yet, been able to demonstrate comparable performance or learning benefits. Similarly we have even less understanding of how (and if) computational

Palmiter et al., 1991; Philips, 1986), thereby preventing them from having the equivalent computational benefits that static diagrams offer.

In sum, we propose a new agenda for research into graphical representations that is based on an analysis of interactivity and, thus, considers the relationship between different external and internal representations. Such an approach should help us to better understand, design and select graphical representations – be they ‘old fashioned’ or technologically advanced – which are appropriate for the learning environment, problem-solving task or entertainment activity in question.

Acknowledgements

The authors would like to gratefully acknowledge support for this work from the ESRC (Cognitive Engineering Programme award L127251033). We should also like to thank the anonymous reviewers of this paper as well as Ben du Boulay, Ron Chrisley, Graham Hole and other members of the Socio-Cognitive Research Group at COGS, Sussex University for their comments on early drafts.

References

- Amigues, R., & Caillot, M. (1990). Les representations graphiques dans l'enseignement et l'apprentissage de l'electricite. *European Journal of Education* 4, 477-488.
- Anderson, J. R. (1990)

372-378.

Cheng, P. C-H. (1993) Multiple interactive structural representation systems for education.

In: R. Cox, M. Petre, P. Brna & J. Lee, Eds. *Proceedings of the International Conference on Artificial Intelligence in Education* (AI ED '93), August, Edinburgh. pp. 25-27.

Cover A. S., Ezquerra, N. F., O'Brien, J. F., Rowe, R., Gadacz, T. & Palm, E. (1993)

Interactively deformable models for surgery simulation. *IEEE Computer Graphics and Applications*, (6), 68-85.

Cox, R. & Brna, P. (1993) Analytical reasoning with external representations. In: R. Cox,

M. Petre, P. Brna & J. Lee, Eds. *Proceedings of the International Conference on Artificial Intelligence in Education* (AI ED '93), August, Edinburgh. pp. 33-36.

Cox, R. & Brna, P. (1995). Supporting the use of external representations in problem solving: The need for flexible learning environments. *Journal of Artificial Intelligence in Education*, (in press).

Earnshaw, R.A. & Watson, D., Eds. (1993) *Artificial Intelligence in Education: Applications* London: Academic Press.

Fairchild, K. M., Serra, L., Hern, N. H., Hai, L.B. & Leong, A.T. (1993) Dynamic FishEye Information Visualizations. In: R. A. Earnshaw, M. A. Gigante and H. Jones, Eds. *Artificial Intelligence in Education: Applications* pp. 161-177. London: Academic Press.

Finke, R.A. (1990) *Critical Dimensions in Instruction* Hillsdale, N.J. : Lawrence Erlbaum.

Gigante, M.A. (1993) Virtual Reality: Definitions, History and Applications. In R. A. Earnshaw, M. A. Gigante and H. Jones, Eds. *Artificial Intelligence in Education: Applications* pp. 3-15. London: Academic Press.

Glenberg, A. M., & Kruley, P. (1992) Pictures and anaphora: Evidence for independent processes. *Journal of Experimental Psychology*

Goettl, B. P., Wickens, C. D., & Kramer, A. F. (1991) Integrated displays and the perception of graphical data. *Ergonomics*, 34, 1047-1063.

Goodman, N. (1968) *Visual Art: An Approach to the Visual Arts*. Indianapolis: Bobbs-Merrill.

Green, T. R. G. (1989) Cognitive dimensions of notations. In A. Sutcliffe & L. Macaulay, Eds. *Open Computers*. Cambridge: Cambridge University Press.

Green, T. R. G. (1990) The cognitive dimension of viscosity: a sticky problem for HCI. In: D. Diaper, D. Gilmore, G. Cockfield, P. Shaker.

Wokingham, England: Addison-Wesley.

Koedinger, K. R., & Anderson, J. R. (1990) Abstract planning and perceptual chunks: elements of expertise in geometry. *Cognition*, 34, 511-550.

Kosslyn, S. M. (1989) Understanding charts and graphs. *Applied Cognitive Psychology*, 9, 185-226.

Kozak, J. J., Hancock, P. A., Arthur, E. J. & Chrysler, S. T. (1993) Transfer training from virtual reality. *Ergonomics* (7), 777-784.

Larkin, J. H. (1989) Display-based problem solving. In: D. Klahr & K. Kotovsky, Eds. *Cognitive Processes in Problem Solving: A Case Study in Chess*. Hillsdale, NJ: Lawrence Erlbaum Associates. pp. 319-341.

Larkin, J. H., & Simon, H. A. (1987) Why a diagram is (sometimes) worth ten thousand words. *Cognition*, 21, (1), 65-100.

Lacience,

n Co un t on August, Edinburgh. pp. 9-16.

Molitor, S., Ballstaedt, S, & Mandl, H. (1989) Problems in Knowledge acquisition from text and pictures. In: H. Mandl & J.R. Levin, Eds. *now A qu s t on ro t n tur s* Amsterdam: North-Holland/Elsevier Science Publishers. pp. 3-30.

Norman, D. (1988) *s o o o Ev r n s* New York: Basic Books.

Norman, D. (1993) Cognition in the head and in the world. *Co n t v n , 1-* (1), 1-6.

- on factual and application learning in an elementary science lesson. *Journal of Computer-Based Instruction* (4), 431-444.
- Rieber, L. P. & Kini, A. S. (1991) Theoretical foundations of instructional applications of computer-generated animated visuals. *Journal of Computer-Based Instruction*, (3) 83-88.
- Scrivener, S. A. R., Harris, D., Clark, S. M., Rockoff, T. & Smyth, M. (1992) Designing at a distance: experiments in remote-synchronous design. In: *Proceedings of the 1992 CHISIG Annual Conference*, Gold Coast Australia, November 1992. pp. 44-53.
- Sheridan, T. B. (1992) Musings on Telepresence and Virtual Presence. *Presence*, (1), 120-25.
- Stenning, K. & Oberlander, J. (1995) A cognitive theory of graphical and linguistic reasoning: logic and implementation. *Cognitive Science*, (19), 97-140.
- Stenning, K. and Tobin, R. (1995) Assigning information to modalities: comparing graphical treatments of the syllogism. In: E. Ejerhed & S. Lindstrom, Eds. *Attention and Cognition* Proceedings of International Conference on Dynamic Semantics, Umea: Sweden. September 1993. Kluwer. (in press)
- Sugiyama, K., & Misue, K. (1991) Visualization of structural information: Automatic drawing of compound digraphs. *IEEE Transactions on Systems, Man, and Cybernetics*, (4), 876-892.
- Sweller, J., Chandler, P., Tierney, P., & Cooper, M. (1990) Cognitive load as a factor in the structuring of technical material. *Journal of Experimental Psychology: Applied*, (2), 176-192.
- Vera, A.H. and Simon, H.A. (1993) Situated Action: A Symbolic Interpretation. *Cognitive Science*, (17), 7-48.
- Winn, W. (1987) Charts, Graphs, and Diagrams in Educational Materials. In: D. M. Willows & H. A. Houghton Eds. *Instructional Design: A Handbook of Research and Applications*. New York: Springer-Verlag. pp. 152-198.

Winn, W. (1993) An account of how readers search for information in diagrams.