



Cognitive Science Approaches To Understanding Diagrammatic Representations

PETER C.-H. CHENG¹, RIC K. LOWE² and MIKE SCAIFE³

¹*ESRC Centre for Research in Development, Instruction and Training, School of Psychology, University of Nottingham, Nottingham, NG7 2RD, UK;* ²*Faculty of Education, Curtin University of Technology, GPO Box U1987, Perth, Western Australia, Australia 6001;* ³*Interactive Media and Virtual Environment Lab, School of Cognitive and Computing Sciences, Sussex University, Brighton, BN1 9QH, UK.*

Abstract. Through a wide variety of approaches cognitive science has given us various important insights into the nature of diagrammatic representations. This paper surveys the findings, issues and approaches to diagrammatic representations in cognitive science. Important current issues that are highlighted include: the relation between the parts of the representational system that are internal to the mind and in external visual media that presents the diagram; the use of multiple representations which is typical of real contexts of diagram use; the benefits of diagrams in terms of (i) computational offloading, (ii) re-representation and (iii) graphical constraining.

Keywords: diagrammatic representations, cognitive science, external cognition, complex information processing

1. Introduction

Cognitive science is a diverse field encompassing many different perspectives for the investigation of a great variety of human cognitive phenomena. The study of reasoning, problem solving and thinking with diagrammatic representations (diagram use) is also diverse, ranging from work on the analysis of the characteristics of diagrams in themselves to studies of mental imagery. This paper reviews research on the nature of diagrammatic representations and what makes them effective, with a particular focus on the issues that are current in the area.

To set the scene consider three diagrams, which will be considered occasionally throughout the paper. We can all recognize that Figure 1 is a weather map for Australia. Most readers will know that the thick contour lines are isobars, or lines of equal pressure. Except for those trained in meteorology, none of us will be able to forecast the weather from the map nor say how the area of pressure will change over time, but such predictions can be made using




Figure 1. A weather map for Australia.

the map. There is clearly more going on here in cognitive terms than can be explained by simply saying that the map is a pictorial image and interpreted as such.

Similarly, consider a secondary school child trying to understand the path of carbon through the environment. Typically this is depicted in textbooks as a cyclical representation, involving text, pictures/schematics and a set of conventional notations such as arrows, lines or boxes (e.g. Figure 2). Children – and indeed adults – often have great problems in understanding this kind of representation at other than a superficial level, despite the inclusion of pictures (icons) and text (cf. weather map). Why is this the case and how can a cognitive science approach improve on the situation?

Figure 3 shows diagrams used by early physicists to discover the conservation of momentum and energy, in the context of head-on collisions of particles moving in a straight line. In each diagram, the labelled lines denote properties of the domain: the initial velocities (U), final velocities (V) and masses (m) of two bodies (subscripts 1 and 2) for a single collision. The orientation and lengths of the U and V lines represent the direction and speed of the bodies, and the relative lengths of the m lines are in proportion to the masses of the bodies. There are obvious fundamental visual differences between these diagrams and the weather map and Carbon cycle diagram, such as their geometric character. But more interestingly, what can be said in cognitive

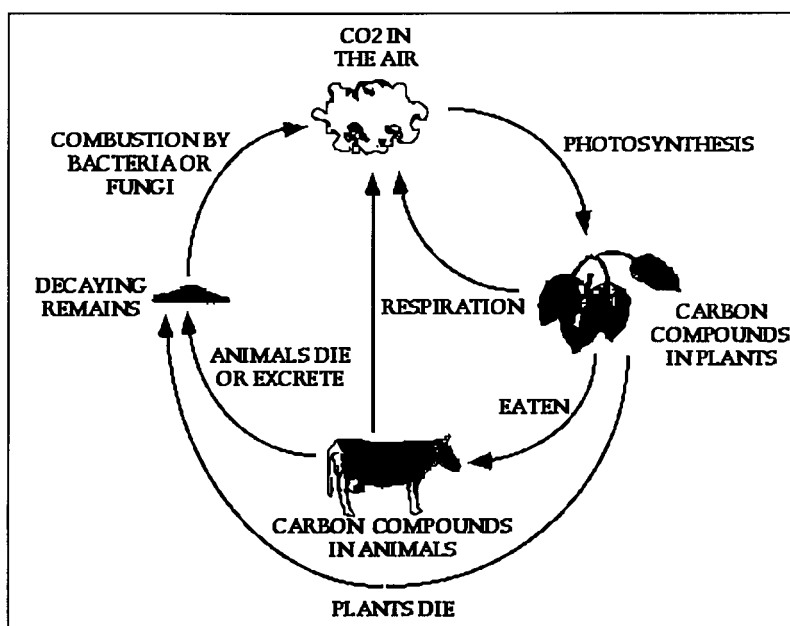


Figure 2. Diagram of the carbon cycle.




Figure 3. Diagrams for elastic collisions.

terms about their similarities that will allow us to understand how diagrams are used and what makes effective diagrams for learning and training?

In three main sections the paper considers (§3) the fundamental nature of diagrams, (§4) aspects of cognition with diagrams, and (§5) what makes diagrams effective. Within each section the issues are addressed from two different but complementary perspectives. The first perspective is a relatively

general one, which considers diagrammatic representations in terms of high level characteristics. The second, at a more specific level, considers the nature of the complex information processing, CIP, that occurs with diagrammatic representations. A brief explanation of how CIP theory applies to diagrammatic representations is given in the next section, before we turn to the main sections.

2. Applying Complex Information Processing Theory To Diagrams

Although new perspectives have developed in cognitive science, the traditional approach that characterizes cognition as complex information processing in terms of physical symbols systems using heuristic search (exemplified by Newell and Simon 1972) is still as relevant and productive today as it has always been. Since the early days of the theory of complex information processing (CIP) systems, in which a narrow range of relatively simple puzzles and games were studied (Newell and Simon 1972), the scope of CIP theory has expanded to cover diverse phenomena, from low level perceptual-motor skill learning, through to design and invention, even collaborative scientific discovery, and of course the use of diagrams.

The heart of the CIP theory is the characterization of human cognition in terms of the procedures that process information in the form of expressions, assemblies of symbols. In diagrams the symbols are visual features, such as the shape, size, orientation of graphic objects. Consider the domain of elastic collisions between two bodies moving in a straight line. Each diagram in Figure 3 is an expression, an arrangement of visual symbols standing for properties of the referent objects involved in the collisions.

Typically, sets of permissible expressions, problem state spaces, are large and organised hierarchically. However, the human information processing system operates in a mainly serial fashion, thus for effective problem solving heuristics are used to guide the search through the state space. Characterizing the nature of the symbol expressions, operators and heuristics for particular tasks or domains is central to the understanding of problem solving. The information processing system searches for expressions, states of knowledge, that will achieve the goals of a given task or problem by using operators or rules to create, modify, reproduce and destroy expressions. For example, to transform the “default” collision diagram, Figure 3a, into a new diagram, Figure 3b, representing the collision between bodies of very different mass, requires operators that modify the lengths of the lines within given constraints. For example, the lengths of m_2 is shortened and m_1 lengthened, whilst keeping the total length of the two lines constant. Other geometrical operators are applied to change length or position of v_1 and v_2 .

An important consequence of the CIP approach is the provision of a criterion on which to base judgements of the relative merits of different representations (e.g., diagrams vs. sentences). This is the notion of *information equivalence*, which recognises two representations as equivalent when all the information (content of expressions) in one is also inferable from the other, and vice versa (Palmer 1977; Larkin and Simon 1987). The application of this criterion permits the investigation of representations with different overall formats but that are at a fundamental level the same. Thus, it is legitimate to attribute any benefits of a representation to its cognitive or computation properties rather than merely because it contained more information in the first place.

Along with general perspective of high level aspects of diagrammatic representations, the CIP perspective will provide a basis for the consideration of the main issues of cognition with diagrams, in the next three sections.

3. Nature of Diagrams

One reason for the persistent interest in diagrams in diverse fields including computer science, education and psychology comes from the common intuition that diagrams are often more effective than propositional representations for whatever purpose they are put. The study of diagrams in cognitive science challenges this intuition and some of the basic distinctions that are held about diagrams.

Diagrams are (sometimes) better

Claims in the literature that diagrams are better, a priori, than other representations should be treated with caution. One finds in surveys of the research in computer science, psychology and education claims about the benefits of diagrams, and visual representations more generally, that seem to be motivated largely by intuitions and that are only weakly supported by rigorous empirical evaluations or any consistent attempt to derive generalisable theories.

For example, consider the case of information representation using multimedia, which seems to offer the possibility of improvement over conventional alternatives in displaying diagrammatic information. There is a wide-spread belief that representations rendered as, say, computer animations have distinct advantages over their paper-based equivalents. However, without any understanding of what makes external representations effective design will continue to be – as it is now – driven by slogans such as ‘a picture is worth 10000 words’, that ‘more is more’ or that the ‘sum is greater than the parts’

(e.g. Lopuck 1996). Such beliefs are, at best, unsupported as a general claim and often seem to rest on unproven and naive assumptions about the way that external visual representations ‘produce’ internal models (Scaife and Rogers 1996). Current orthodoxies about the intrinsic benefits of visualisation of information, which are grounded on the assumption that it makes the information more accessible, need to be examined far more critically.

It is noteworthy that the cognitive scientists Larkin and Simon (1987) included a qualification (in parentheses) in the title of their seminal paper, which will be considered below – “Why a diagram is (sometimes) worth ten thousand words”. As will be seen below, a cognitive science perspective on the use of diagrams reveals a highly complex phenomenon with many facets, from which it is not possible to simply derive straightforward general claims about the benefits of diagrams.

Diagrams and propositional representations may not be so different

One reason why it is not possible to claim that diagrams are generally better than other representations is that they are not so different from other representations. A distinction is often made between diagrammatic and propositional representations, such as logic and mathematics. The main characteristics of diagrams is their use of space and spatial properties (location, topology, geometry, etc), but this is not exclusive to diagrams, because propositions also use “diagrammatic” properties to encode information, although to a lesser degree (e.g., in the formula ‘ $x = y + z$ ’, it matters whether the ‘+z’ term is to the left or right of the equals sign). By the same token diagrams are not purely diagrammatic, because they contain propositions, as in the Carbon cycle diagram (Figure 2). One may consider all representations as falling at different places in a continuum from little use of diagrammatic properties to encode information through to substantial use of such properties. Thus, strong claims about the difference between diagrams and other representations should be treated with caution.

Considering the nature of the information that is being processed there is an *a priori* reason to treat diagrams as a distinct class of representational systems. Under the CIP approach in cognitive science, the assumption is that differences between representational systems can best be understood in terms of their respective symbol expressions, operators and heuristics. In this fashion, Larkin and Simon’s (1987) paper compared diagrams with equivalent sentential representations, explaining that the cognitive benefits of many diagrams reside in the way that information, symbol expressions, are indexed by spatial location rather than by symbolic labels.

Diagrams are not a unitary class

Another reason why diagrams cannot be assumed to be generally better than some other class of representation is that diagrams are not a well-defined unitary class. There are many different types of diagrams, for example, just compare the map, flow/cycle diagram and the geometric diagram in this paper (Figures 1, 2, 3). The claims in the literature about diagrams in general, which are derived from studies that examined just a single type of diagram should be treated with caution.

Again, by considering the nature of the information that is being processed, the variety of diagrammatic representational systems are more clearly distinguished, because of the focus on symbol expressions, operators and heuristics. For example, compare the diagrams above. In Figure 3, lines for properties of bodies (symbols) are related by the geometric structure of the diagram (expressions), and a diagram may be modified using geometric rules (operators). In Figure 2, icons stand for the location of CO₂ in different entities and the labelled lines represent processes (symbols), and the combination of the arrows between two icons shows the transfer of CO₂ (expression). The diagram might be modified to include fossilization and the burning of hydrocarbon fuels, by adding more icons and connecting them together with appropriately labeled arrows (operators). Similar, analysis of Figure 1 is left for the reader. Such comparisons demonstrate the huge variety of diagram types, arguably more diverse than propositional representations.

4. Aspects of Cognition With Diagrams

Cognitive science has taken a number of perspectives when studying the use of diagrams, which focus on different aspects of the relation between the diagram user and the nature of the diagram itself.

At a fundamental level, a diagrammatic display can be regarded as an arrangement of various graphic elements in space. Perceptual similarities and differences between these elements allow them to be grouped or distinguished according to visuospatial characteristics of the particular display. For example, in the weather map diagram shown in Figure 1, we could group together the series of concentric curves in the northeast because of their similarity in shape, graphic treatment and location. These are readily distinguished from the bold lines bearing triangular barbs in the southeast of the diagram. Being able to configure diagram elements into groups or discriminate between them in this way is an important precondition for proper interpretation and is a key requirement of a well-designed diagram.

An explanation for why this is so comes from the consideration of the nature of information processing involved. Larkin and Simon (1987) suggested that it is the use of locational indexing in diagrams that often makes them more effective than informationally equivalent sentential representations. This form of indexing means that information that tends to be needed for the same inference can usually be found in adjacent locations in a diagram, so reducing the amount of search required to find the information. Further, perceptual inferences with diagrams allow the power of the highly parallel human visual system to replace more cumbersome serial logical inferences. Such inferences do not pose a fundamental problem for diagrams under the CIP approach, as they may be treated as operators. Seeing that the lengths of the lines in Figure 3a are equal is like testing for the equality of the values assigned variables in a mathematical representation. Studies of geometry problem solving (Koedinger and Anderson 1990) and reasoning with electrical circuit diagrams (Egan and Schwartz 1979), for example, illustrate how perceptual inferences can be dealt with under the CIP approach in cognitive science.

The raw perceptual information that a diagram provides for the user by way of visuospatial cues must be modulated by knowledge about the individual and collective meanings of the graphic elements. An over-reliance on the visuospatial characteristics of the markings making up a diagram can be highly misleading. This can be demonstrated by two other sets of elements shown in Figure 1. The sets of roughly concentric markings in the southwest and southeast corners of the diagram respectively are widely separated and so, in purely perceptual terms, appear to be quite distinct. When beginning students of meteorology (novices) were asked to group the elements on this diagram, they typically distinguished between the sets of markings in these two corners of the diagram (Lowe 1993). In contrast, professional meteorologists (experts) configured these two sets as a single grouping. It seems that a major factor determining a viewer's capacity to make effective use of a diagram is how much that person already knows about the sort of subject matter depicted in the diagram and the specific method of depiction.

Using a given diagram effectively requires the viewer to think about that diagram in quite particular ways. Because diagrams are highly specialised depictions that differ substantially from more realistic pictures, the cognitive approaches that we habitually use for interpreting our everyday visual environment are inappropriate. The skills required for using diagrams effectively must be learned and appear to be highly domain-specific. There are some generic aspects that influence diagram use (such as the need to treat these as abstract rather than literal depictions).

An implication that can be drawn from this research is that diagrams, in and of themselves, do not ‘contain’ all the information that a viewer needs to use them properly. Rather, the background knowledge that the viewer brings to the diagram plays a critical role in whether or not it can be processed satisfactorily. This would mean that good diagram design can only go so far in determining whether a diagram is likely to be an effective way for depicting particular information. For this reason, current orthodoxies about the intrinsic benefits of visualization of information (on the assumption that it makes the information more accessible) need to be examined far more critically.

The discussion of Figure 1 has so far focused upon it as a representation of a particular state and shown the importance of domain-specific knowledge in effective diagram use. However, diagrams are frequently used in more sophisticated ways that involve mental processes such as inference and prediction. For example, a weather map diagram for a particular day can be used to make a prediction about the weather pattern that is expected on the following day. In this case the diagram is the basis for generating new information rather than simply depicting the present situation. Similarly, by modifying Figure 3a for symmetrical collision between bodies, one may, for instance, explore possible asymmetrical configurations, such as Figure 3b. The cognitive processes involved in this type of task require the creation of a suitable mental model that can be ‘run’ to make predictions or inferences, in the case of the weather map this may even be backwards as well as forwards in time.

Similarly, multimedia design provides a strong challenge for any general theory of external representations but also emphasises these issues which are central to understanding cognition with diagrams. Firstly there is the possibility that being able to *interact* with multimedia representations in ways not possible with single media (i.e. books, audio, video), can lead to easier learning, better understanding, and increased motivation. This is certainly the case but leads to the question of how users interact with any kind of diagrammatic representation. This is often overlooked in studies of paper-based diagrams but is surely important. For example marking the paper or making other annotations are a central feature of geometry student progress (e.g. Koedinger and Anderson 1990). Thus we need to recognise the important role of *constructing* external representations (Reisberg 1987), which is normally such an integral part of learning or problem-solving, e.g. underlining, making notes separately, re-representing text-based ideas in various diagrammatic forms, sketching etc.

Multimedia also affords novel access to *multiple* representations of information. An example is multimedia encyclopaedias which have been designed on this principle, providing a variety of audio and visual materials on any

given topic. However, the issue of the benefits of multiple representations is also present for paper-based products – consider even the embedding of pictures within text. To produce effective designs it is necessary to understand how learners integrate information arising from different representations of the same and different information. This requires analysing how people learn to *read* and comprehend the significance of the content the diagram, for example how they develop an understanding of canonical diagram forms, and how this is assimilated to their current understanding of the domain.

These kinds of issues underpin the need for a more general account of diagrams, *qua* external representations, than the case-based approach which has dominated the research literature. There have been a number of approaches, albeit different, that seem highly promising in this regard. One is the work of Stenning and colleagues on paper- and computer-based representations (e.g. Stenning and Oberlander 1995; Stenning and Lemon 2001; Stenning 1999) who argue for the need to distinguish between ‘expressive’ and ‘processing’ explanations for the cognitive usability of diagrams. The former has to do with semantic constraints on the space of diagrammatic interpretations, the latter to do with perceptual and/or mnemonic limitations due to the way that the particular diagram (or other representational form) is constructed. Another approach is that of Zhang and colleagues (e.g. Zhang and Norman 1994; Zhang 1997) who emphasise the mappings between rules and the structure of the problem space, both internal and external. Finally there is the work of Green and colleagues (e.g. Green and Petre 1996) who stress the value of high-level abstractions to convey important characteristics of external representations, such as the complex interactions between parts of the representational system that lead to ‘viscosity’ – a resistance of any part of the representation to local change.

Scaife and Rogers have an approach they label ‘external cognition’, that focuses on how different representations are processed when performing different activities (Scaife and Rogers 1996). The emphasis here is on the interactions between internal and external representations considered together (cf. Larkin 1989; Norman 1993; Vera and Simon 1993). Their belief is that the process by which different external representations are used in learning or problem-solving is complex, involving an interaction between internal processes and different aspects of external representations at different stages of a task. For example, reading and abstracting knowledge from a diagram requires making connections between different elements of the display in a temporal sequence. Such a ‘take’ may be contrasted with accounts that either emphasise the primacy of internal representations and/or ignore the way they are co-ordinated with external ones.

At the information processing level the same set of issues is cast in terms of the relation between the aspects of the representation that are internal to the mind and those that are in the external environment. For example, Figure 3 is an external physical notation but the geometric rules to manipulate the diagrams are usually held in the user's memory. All but the most trivial problems require iterative cycles of (i) visual interpretation of the external diagrams, (ii) internal recognition of applicable operators, (iii) modification of the drawing, and (iv) further visual interpretation of the new diagrams. Larkin (1989) and Zhang and Norman (1986) consider how the distribution of representations between the mind and external environment may reduce working memory loads and lessen cognitive demands. Tabachneck and Simon (1994) present a model of how internal images, in the "mind's eye", could be processed.

Under the CIP view there is a recognition that both the users and uses of diagrams should be considered. Users of a notation in a particular domain who are, for example, more expert will engage different operators and heuristics. In effect they possess quite a different representational system to novices in the same domain, even though they may share a common external notation. In the same vein, different tasks have distinct goals, which will be satisfied by different information. The search for goal expressions may require alternate operators and heuristics to process the notation, and may even be considered to constitute different representations, in some cases. Cheng (1996) discusses some of the variety of tasks or *functional roles* that diagrammatic representations may support.

If so much of the capacity to use a diagram effectively is bound up with what the viewer already knows about the subject matter, what options are available for improving diagram use? This question should be of particular interest to educators who provide novices in a domain with diagrams on the assumption that they will make the subject matter more accessible. The problem seems to be one of 'boot-strapping'; without a certain minimum knowledge of the domain, an individual is unlikely to be able to use a domain-specific diagram effectively. One approach for addressing this issue is to help novices develop the sorts of basic knowledge structures that could support appropriate cognitive processes. This type of approach has been explored recently by providing meteorological novices with computer-based animations designed to act as external models that could help them to build mental models of weather map systems that are more consistent with those used by experts in the field (Lowe 1997).

As mentioned with respect to Figure 1, when beginning students of meteorology (novices) were asked to group the elements on this diagram, they typically distinguished between the sets of markings in these two corners of

the diagram (Lowe 1993). In contrast, professional meteorologists (experts) configured these two sets as a single grouping. Further investigations indicated that the experts' knowledge of the wider context of Australian weather systems allowed them to relate these sets of markings meteorologically as two sections of a much larger-scale feature that connected them beyond the scope of the diagram (Lowe 1994). In contrast, the novices were unable to invoke this type of domain-specific knowledge and so appeared to be reliant solely on visuospatial information (Lowe 1996).

Comparisons of meteorological experts and novices suggest that the superior quality of experts' predictions of weather map patterns is related to particular characteristics of the mental model they construct from a given weather map (Lowe, in press). Not only do they appear to be able to construct more extended and detailed mental models of the depicted situation, they also rely on a rich store of knowledge about the properties and behaviour of the various meteorological features. Once again, a key factor in using a diagram effectively is what the viewer brings to the diagram (rather than what the diagram brings to the user).

This is consistent with Koedinger and Anderson's (1990) work on the differences between novice and expert geometry problem solvers. They discovered that the problem solvers search a space of perceptual chunks comprising meaningful diagrammatic configurations, so performed better than novices who deal with the visual elements of the same diagrams in a piecemeal fashion.

5. Properties of Effective Diagrams

In principle, at a basic level, it is obvious that a well-designed diagram should allow the user to make a relatively straightforward mapping between the diagrammatic depiction and the situation it represents. This means that it should be a simple matter to compare each component in the represented situation with its corresponding component in the diagram ("This is Australia"). It should also be easy to compare the corresponding arrangements of these components between situation and diagram ("These are concentric isobars over Australia").

However, given the richness of the nature of cognition with diagrams discussed in the previous section, considerations of what makes a diagram more or less useful must take a broader view. Clearly the effective properties will vary with the particular diagram and situation of use but Rogers and Scaife (1999) identify at least the following kinds of 'computational offloading' – the ways in which different external representations reduce the amount of cognitive effort required to solve informationally-equivalent

problems (e.g. Larkin and Simon 1987): (i) Re-representation – This refers to how different external representations, that have the same abstract structure, make problem-solving easier or more difficult and how they are selected (e.g. Zhang and Norman 1994); (ii) Graphical constraining – This refers to the way elements in a graphical representation are able to limit the range of inferences that can be made about the represented concept (e.g. Stenning and Oberlander 1995); (iii) Temporal and spatial constraining – This refers to the way different representations can make relevant aspects of processes and events more salient when distributed over time and space (e.g. the use of canonical cyclical diagrams, as in the carbon cycle of Figure 2).

Clearly other properties could be identified as the advantages of diagrammatic over other kinds of representation but a major task remains that of understanding how these properties are actually realised, which will necessitate a better understanding of the mechanisms relating internal and external representations.

At the information processing level, the same issues can be addressed, but it is first necessary to consider what the appropriate bases are for the comparison of different representational systems. For representations that are informationally equivalent (see above), Larkin and Simon (1987) demonstrated that diagrams often have computational advantages over sentential representations, as already noted. However, if diagrams that are not informationally equivalent are to be studied, the information processing approach provides other bases for comparisons of representations. At a low level, comparison can be made between the form, number and complexity of the symbols, expressions and operators of different representational systems. At a higher level, comparisons can be made in terms of the overall size and/or complexity (e.g., breadth and depth) of the problem state space for the representations. For example, Cheng and Simon (1992) showed that in the inductive discovery of the law of momentum conservation the overall space of expressions is smaller for a diagrammatic representation (similar to Figure 3) than it is for an algebraic notation.

6. Conclusion

This review has examined cognitive science approaches to understanding diagrammatic representations. Below the surface of common but somewhat naive claims about the benefits of diagram over other representations lie various complex cognitive issues that inform us about the nature of human understanding, problem solving and thinking more generally. Diagrams are sometimes, perhaps often, better than other representations, but the reasons

why are complex. To conclude, a summary of the main issues covered at various points throughout the paper is presented.

- 1) Claims in the literature that diagrams are better, a priori, than other representations with respect to presenting information should be treated with caution.
- 2) Diagrams are not a unitary class of representations but (i) are similar in some important respects to propositional representations and (ii) come in a wide variety of forms which may have quite different implications for cognition.
- 3) Properties that make diagrams effective are shared by many other representations.
- 4) There are diverse uses for diagrams which may have quite different implications for cognition.
- 5) Diagrams are hardly ever found in isolation, so the way that multiple representations are simultaneously used for reasoning and learning is an important issue.
- 6) The study of diagram use should examine the cognitive processes involved in diagram interpretation and understanding and not just the perceptual properties of graphic displays.
- 7) Similarly, the interactive processes of diagram construction and modification should be considered in addition to the interpretation of diagrams.
- 8) There are internal and external aspects of diagram use that need to be explained, including the role of background knowledge and the role of diagrammatic conventions – learning to recognise canonical forms.
- 9) The contrast between expert and novice users of diagrams is an effective way to learn about what makes diagrams effective or not.
- 10) Some of the properties that can (sometimes) make diagrams particularly effective representations have been identified in terms of their effectiveness in promoting computational offloading, for example:
 - i) The locational indexing of information.
 - ii) Re-representation by selection of more powerful operators or redistribution of the internal/external distribution of the elements of the representations.
 - iii) Graphical constraining in limiting the size and complexity of the search space.
 - iv) Temporal and spatial constraining making processes and events more salient when distributed over time and space.

References

- Cheng, P.C.-H. (1996). Functional Roles for the Cognitive Analysis of Diagrams in Problem Solving. In Cottrell, G.W. (ed.) *Proceeding of the Eighteenth Annual Conference of the Cognitive Science Society*, 207–212. Lawrence Erlbaum: Hillsdale, NJ.
- Cheng, P.C.-H. & Simon, H.A. (1995). Scientific Discovery and Creative Reasoning with Diagrams. In Smith, S., Ward, T. & Finke, R. (eds.) *The Creative Cognition Approach*, 205–228. MIT Press: Cambridge, MA.
- Egan, D.E. & Schwartz, B.J. (1979). Chunking in Recall of Symbolic Drawings. *Memory and Cognition* 7(2): 149–158.
- Green, T.G.R. & Petre, M. (1996). Usability Analysis Of Visual Programming Environments – A Cognitive Dimensions Framework. *Journal of Visual Languages And Computing* 7: 131–174.
- Koedinger, K.R. & Anderson, J.R. (1990). Abstract Planning and Perceptual Chunks: Elements of Expertise in Geometry. *Cognitive Science* 14: 511–550.
- Larkin, J.H. (1989). Display-Based Problem Solving. In Klahr, D. & Kotovsky, K. (eds.) *Complex Information Processing: The Impact of Herbert A. Simon* 319–341. Lawrence Erlbaum Associates: Hillsdale, New Jersey.
- Larkin, J.H. & Simon, H.A. (1987). Why a Diagram is (Sometimes) Worth Ten Thousand Words. *Cognitive Science* 11(1): 65–100.
- Lopuck, L. (1996). *Designing Multimedia*. Peachpit Press: Berkley, CA.
- Lowe, R.K. (1993). Constructing a Mental Representation from an Abstract Technical Diagram. *Learning and Instruction* 3(3): 157–179.
- Lowe, R.K. (1994). Selectivity in Diagrams: Reading Beyond the Lines. *Educational Psychology* 14: 467–491.
- Lowe, R.K. (1996). Background Knowledge and the Construction of a Situational Representation from a Diagram. *European Journal of Psychology of Education* 11: 377–397.
- Lowe, R.K. (1997, August). Effects of Interactive Animation in Complex Visual Learning. *Paper Presented at the Seventh European Conference for Research on Learning and Instruction*. Athens, Greece.
- Lowe, R.K. (in press). Domain-Specific Constraints on Conceptual Change in Knowledge Acquisition from Diagram. In Vosniadou, S., Carretero, P. & Schnotz, W. (eds.) *New Perspectives on Conceptual Change*. Erlbaum.
- Lopuck, L. (1996). *Designing Multimedia*. Peachpit Press: Berkley, CA.
- Newell, A. & Simon, H.A. (1972). *Human Problem Solving*. Prentice-Hall: Englewood Cliffs, NJ.
- Norman, D. (1993). Cognition in the Head and in the World. *Cognitive Science* 17(1): 1–6.
- Palmer, S.E. (1977). Hierarchical Structure in Perceptual Representation. *Cognitive Psychology* 9: 441–474.
- Reisberg, D. (1987). External Representations and the Advantages of Externalizing One's Thoughts. *Proceedings of the IX Annual Conference of the Cognitive Science Society*. Seattle, WA/Hillsdale, NJ: LEA.
- Rogers, Y. & Scaife, M. (1999). How can Interactive Multimedia Facilitate Learning? *Proceedings of First International Workshop on Intelligence and Multimodality in Multimedia Interfaces*. Menlo Park, Ca: AAAI Press (in press).
- Scaife, M. & Rogers, Y. (1996). External Cognition: how Do Graphical Representations Work? *International Journal of Human-Computer Studies* 45: 185–213.
- Stenning, K. (1999). Distinguishing Semantic from Processing Explanations of the Usability of Representations: Applying Expressiveness Analysis to Animation. *Proceedings of*

- First International Workshop on Intelligence And Multimodality in Multimedia Interfaces*. Menlo Park, CA: AAAI Press (in press).
- Stenning, K. & Lemon, O. (1997). Diagrams and Human Reasoning: Aligning Logical and Psychological Perspectives. Position paper for 'Thinking With Diagrams II' Conference, University of Portsmouth, January.
- Stenning, K. & Oberlander, J. (1995). A Cognitive Theory of Graphical and Linguistic Reasoning: Logic and Implementation, *Cognitive Science* **19**: 97–140.
- Tabachneck, H.J.M., Leonardo, A.M. & Simon, H.A. (1994). How Does an Expert Use a Graph? A Model of Visual and Verbal Inferencing in Economics. In Ram, A. & Eiselt, K. (eds), *Proceedings of the 16th Annual Conference of the Cognitive Science Society*, 842–847. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Vera, A.H. & Simon, H.A. (1993). Situated Action: A Symbolic Interpretation. *Cognitive Science* **17**(1): 7–48.
- Zhang, J. (1997). The Nature of External Representations in Problem Solving. *Cognitive Science* **21**(2): 179–217.
- Zhang, J. & Norman, D. A. (1994). Representations in Distributed Cognition Tasks. *Cognitive Science* **18**(1): 87–122.