PERGAMON

Computers & Education 33 (1999) 109–130

COMPUTERS & EDUCATION

www.elsevier.com/locate/compedu

Unlocking conceptual learning in mathematics and science with effective representational systems

Peter C.-H. Cheng

ESRC Centre for Research in Development, Instruction & Training (CREDIT), University of Nottingham, Nottingham NG7 2RD, UK

Abstract

The representational analysis and design project is investigating the critical role that representations have on conceptual learning in complex scientific and mathematical domains. The fundamental ideas are that the representations used for learning can substantially determine what is learnt and how easily this occurs, and that to improve conceptual learning effective representations should be found or invented. Through the conceptual analysis and empirical evaluation of a class of representations that appear to be particularly beneficial for conceptual learning, Law Encoding Diagrams (LEDs), the project has identified certain general characteristics of effective representations. In this paper a descriptive model of the components and processes of conceptual learning is presented and used for several purposes: to explain why the nature of representation used for learning is critical; to demonstrate how representations possessing the identified characteristics of effective representations appear to support the major processes of conceptual learning; to consider how computers may further enhance the potential benefit of LEDs for conceptual learning. © 2000 Elsevier Science Ltd. All rights reserved.

1. Introduction

One of the major challenges now facing the use of computers in education is to develop systems that effectively support conceptual learning in substantial mathematical and scientific domains. Computer-based learning has had some success at promoting procedural or skills-based learning (e.g. Anderson, Corbett, Koedinger, & Pelletier, 1995; Lesgold, Lajoie, Bunzo, & Eggan, 1992; see also papers by Wood and Wood, and Wood, Underwood and Avis in this volume), but it is by no means obvious that the design of these systems can be extended to support conceptual learning.

This paper summarizes the theoretical and empirical aspects of a programme of research that

PII: S0360-1315(99)00028-7

is investigating how best to support conceptual learning. The basic claim of this work is that the nature of the representational systems used for such learning, to a significant extent, determine what is learnt and how easy it is to obtain a good conceptual understanding. The keys to conceptual learning are effective representations; the programme of research will be referred to as 'the representational analysis and design project'.

Work in cognitive science on problem solving, expertise and the nature of representations has demonstrated the pivotal role that representations have in human cognition, especially in reasoning and problem solving. Following Skemp (1971), Kaput (1992) and others, the representational analysis and design project is extending this work into the arena of learning and in particular to conceptual learning. Motivation for the approach of the present project comes, in part, from studies of the nature of scientific discovery which have demonstrated that the role of representations are fundamental to the success of the scientific enterprise. Major figures in the history of science made their discoveries by carefully selected representations, or by inventing new representations. These representations had informational and computational properties that facilitated the discoveries. Thus, an appealing notion is that learners may also benefit substantially in similar ways by being given carefully selected representations, or even specially invented representations, for use in conceptual learning.

To introduce some of the issues to be covered in this paper, consider two learners' attempts at solving a particular physics problem: one using a conventional algebraic approach and the

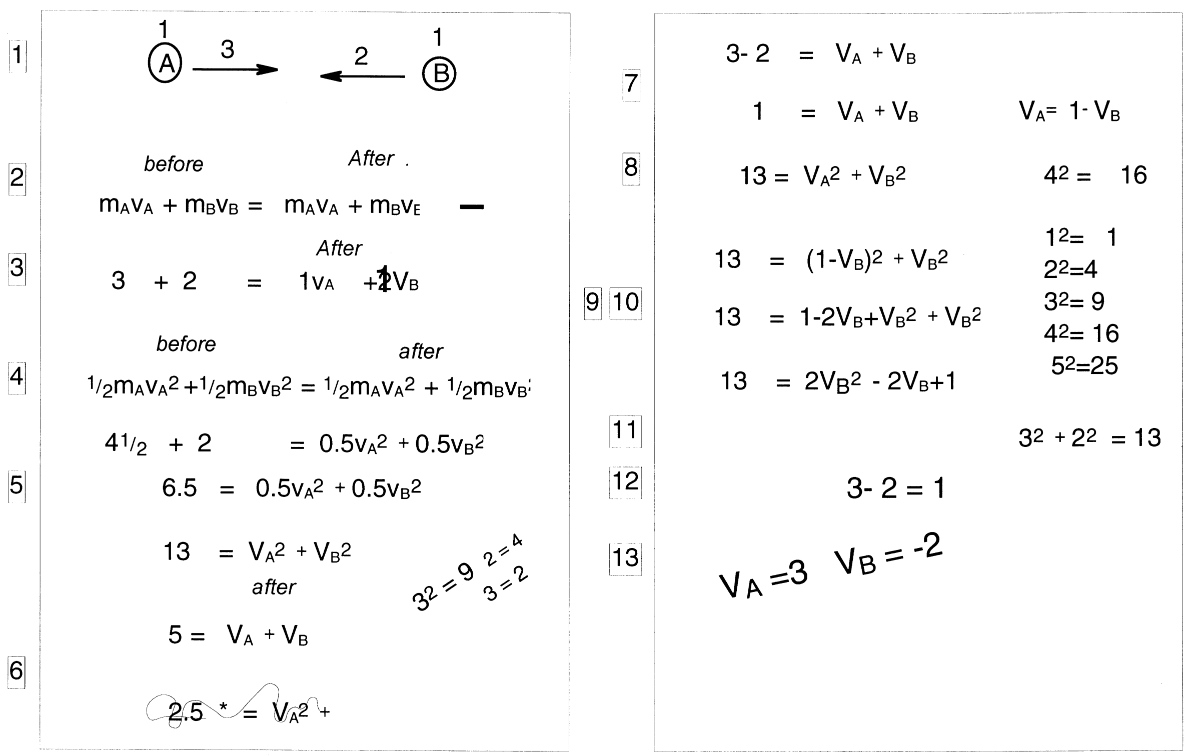

Fig. 1. PB's algebraic solution to a particle collision problem.

second an approach using a diagrammatic representation. Both solutions were produced by participants in the empirical evaluations that have been conducted on the effectiveness of different representations for learning as part of the present project. The problem involves finding the velocities of two elastic bodies which have collided head-on, given the masses of the bodies and their initial velocities.

A picture illustrating this class of problems is shown as the first step of the solution attempt in Fig. 1. (The two pages have been accurately redrawn for clarity.) This solution was produced by experimental participant PB, who was a graduate physicist working for his doctorate in Physics. This attempt was generated in the pre-test of an experiment. PB would have been well schooled with the relevant laws and algebra during his undergraduate degree. Given the initial state of the problem, in which body-A impacts body-B coming in the opposite direction with a different speed, PB drew the diagram showing the initial situation, Step 1 in Fig. 1. The solution involves writing down the algebraic laws for momentum and energy conservation as applied to this class of phenomena, Steps 2 and 4, and substituting in the given values, Steps 3 and 5. At Step 6, PB begins to think about substituting one equation into the other to obtain an equation with just one unknown variable, but he realizes that he has made a mistake in the substitution of values into the momentum equation (at Step 3), so he corrects this before proceeding (Step 7). The full substitution and elimination procedure follows, Steps 8 and 9. PB obtains a quadratic equation but does not solve it. Instead, he reverts to the earlier equation derived from the energy conservation laws (at Step 8) and simply finds a numeric solution by considering squares of integers (Steps 10 and 11). The equation derived from the momentum law (at Step 7) is then used to successfully check this numeric solution (Step 12). PB is happy to state the final solution (Step 13). However, he does not realize that the values are the same as the initial velocity values given in the problem statement, which he had earlier used to label the diagram at Step 1.

Despite being a graduate physicist it is clearly not the case that PB has a good understanding of this basic topic of physics. Unfortunately, PB's performance is not unusual

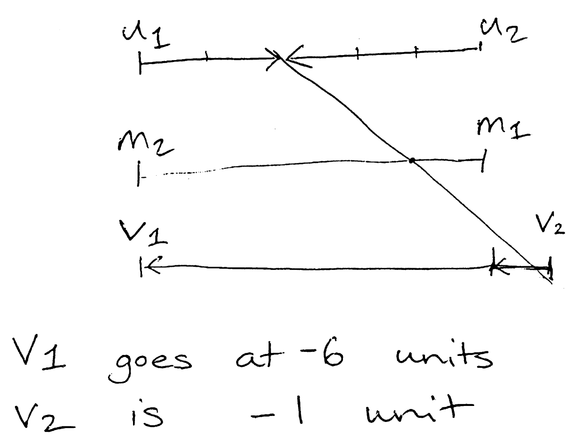

Fig. 2. JT's diagrammatic solution to a particle collision problem.

compared to the other graduate physicists and engineers that have been studied doing similar problems (Cheng, 1996c). This solution attempt is a good example of the difficulties that can arise from the use of a poor representation. Mistakes have been made in the manipulation and interpretation of the expressions generated during the solution, some of which PB has done well to spot and correct, others he has missed completely. Effort has been wasted pursuing a series of inferences in a particular direction which was then abandoned. Knowledge about techniques for the manipulation of the representation itself is necessary for problem solving but is not directly relevant to understanding the nature of the domain, in this case methods to solve pairs of simultaneous equations and the solution to quadratic equations. It is easy to lose track of the goal during problem solving when so much algorithmic work has to the done, as shown by PB not realizing that the solution values were identical to the given values.

One of the claims of the representational analysis and design project is that such difficulties can be avoided by giving learners better representations. Fig. 2 shows an alternative approach to the solution using a similar, but more difficult, problem. It is more difficult because the masses of the bodies are not equal. Fig. 2 is an example of a class of representations that was invented by the scientists, Huygens and Wren, who discovered the laws governing this domain (Cheng & Simon, 1995). The magnitude of the masses are shown by the lengths of lines m_1 and m_2 in the middle of Fig. 2. The given initial velocities are shown by the arrows at the top of Fig. 2, u_1 and u_2. The final velocities have been found by construction of the diagram according to certain geometric rules that encode the laws of momentum and energy conservation for this domain. After impact, the bodies depart in the same direction but v_1 is much greater because its mass is much smaller.

The diagram in Fig. 2 was actually produced in the post-test of an experiment, by participant JT who was a graduate psychologist with little knowledge of physics (in contrast to PB, above). JT learnt about the domain for 40 min using a computer-based learning environment that exploits these diagrams and was developed as part of the present project. No direct instruction on the structure or use of the diagrams was given to JT who simply learnt about the domain by manipulating the diagrams on screen and comparing them with an animated simulation of the collisions. Despite the limited amount of instruction and the relatively short time on the system, JT rapidly acquired the rules governing the structure of the diagrams. The post-test was the first time she drew the diagrams for herself, but she successfully applied them to problems she could only guess at before using the system. The experiments conducted in the present project, with a variety of representations in computer-based and conventional media, show that such dramatic changes in approaches to problem solving with new representations is not unusual.

The diagram in Fig. 2 is clearly simpler than the numerous complex expressions in Fig. 1, so the chances of making errors when generating or modifying such representations are likely to be lower. The form of the diagram also reflects the structure of the domain, the pairs of arrows for initial and final velocities giving a simple image of what is happening, so the interpretation of the representation is supported naturally. The diagram uses simple geometric constraints to encode the laws of the domain that are directly applied to elements representing particular properties, whereas the syntactic algebraic rules operate on quite arbitrary expressions remote from the phenomenon itself.

The contrast between Figs. 1 and 2 exemplifies how the representation a learner uses can determine what and how easily they learn. Thus, an approach which appears to have great potential for improving conceptual learning is to judiciously select, or invent, appropriate and effective representations. This paper will summarize the theoretical and empirical work conducted to articulate and test this claim. The paper has three main sections.

The first section considers why it is that representations can exert such an influence on conceptual learning, or to put it another way, to explain why the processes of conceptual learning are so sensitive to the choice of representations. To achieve this, it is necessary to consider the nature of conceptual understanding and the processes of conceptual learning. This is done by presenting a general descriptive model of the components and processes of such learning. This model is used to identify the major barriers that an ineffective representation may raise in the path of the learner; to show that previous computer-based approaches to promote conceptual learning have been largely piecemeal in their attempts to support learners overcoming such barriers, and to argue that the choice of an effective representation can, in a systemic fashion, minimize these barriers.

The second section of the paper also uses this model to provide a context in which to consider the nature of effective representations. Various characteristics of effective representations are identified and described. A particular class of representations, Law Encoding Diagrams (LEDs — illustrated in Fig. 2), has been the main focus of the empirical studies and conceptual analysis in the project, because they possess these characteristics and appear to be effective in supporting conceptual learning. Empirical evaluations of these LEDs have shown that they can improve learning compared with traditional algebraic approaches (Cheng, 1996d). Other systems of LEDs have also been specially designed for various domains (Cheng, 1999a,b,c) and further empirical evaluations have demonstrated that they can promote learning (Cheng, 1999e). An example with LEDs for electricity is presented later.

The third and final section of the paper considers general educational issues and implications concerning the introduction of novel representations.

2. Conceptual learning and representations

The representational analysis and design project is based on a considered choice to explore learning in substantial complex mathematical and scientific domains. Examples are probability theory, electricity, particle collisions, Galilean kinematics, Newtonian dynamics, algebra, differential calculus and thermodynamics, of which the first four have been studied in detail to date. Such domains are being considered because the potential impact that the choice of representations can have on problem solving and learning is much greater than for simpler domains. Important aspects of representational design may be missed in studies with simpler domains, and approaches to support conceptual learning developed for simple domains may not scale up to larger, and more complex ones.

2.1. Conceptual understanding

A firm conceptual understanding of a substantial scientific or mathematical domain will need to encompass all the different ontological, functional and structural aspects of that domain.

Clearly, for a representation to enhance learning effectively, it will need to support each of these to some extent.

For example, consider the domain of electricity: This topic is difficult to learn because it has intangible entities and properties (e.g. electrons, resistance) and uses formal technical concepts (e.g. perfect conductor, Amps). Moreover, the domain involves a great many different types of components (e.g. resistors, wires, switches, bulbs, diodes, etc.) which may be arranged in a huge variety of network configurations (e.g. simple series and parallel arrangements, Wheatstone bridge, to mention a few). The behaviour of these components and networks may be described in concrete terms, by giving values of electrical properties, or they may be characterized in more abstract terms in relation to the general laws of electricity. These laws reflect the complexity of the domain with relations for individual components and also assemblies of components. It is possible to reason about the domain in formal algebraic terms using laws and also in more concrete causal or material terms, perhaps by analogy to hydraulic systems. At a more detailed local level, it is possible to consider the structure of a network in terms of the connectivity of components, topology, or in terms of the spatial arrangement of components.

Assuming that an individual's conceptual understanding of a domain consists of a complex network of concepts (Hiebert & Carpenter, 1992), we may address what constitutes good knowledge by considering the nature of such networks. Precisely what form the internal mental representation takes is not essential for the present considerations (but see Kyllonen & Shute, 1989, and the paper by Gobet and Wood in this volume for discussion of mental knowledge representation). Of concern here is the general structural character of such networks, in relation to the above description of the nature of complex domains.

Some insight into this issue comes from studies of how scientists originally made discoveries in new domains. Contrary to popular stories about discoveries as leaps of insight and serendipity, realistic historical accounts that, for instance, examine the evolution of scientists' ideas through their laboratory notebooks, show that substantial work on the part of scientists is involved. They sift through large bodies of empirical evidence and a host of potential models before they eventually obtain a coherent integration of laws for the domain with the many different manifestations of the phenomena, over a range of complementary but alternative perspectives. If learners of science obtain personal networks of concepts with a similar scope and structure, it would be reasonable to claim that they have a good understanding.

Some of the desirable properties of good networks of concepts are highlighted by Thagard's (1989) Theory of Explanatory Coherence, and its implementation in the ECHO model. The theory and model are used to assess the acceptability of different sides of major debates in the history of science, which demonstrate the possibility that the structure of networks of concepts can be assessed, on a relatively large scale, to determine the quality of understanding. Chi (1992) also considers the ways in which networks of concepts are structured, but she focuses on different hierarchical categories of concepts that may be used to describe aspects of a domain. This theory seeks to explain different degrees of difficulty in achieving conceptual change in the history of science, and in individuals learning science, in terms of different classes of structural changes to the hierarchies of concepts. Again, the point is made that it is feasible and important to address the general nature and structure of networks of concepts when considering the quality of understanding.

Other work in artificial intelligence and cognitive science on models of induction, analogy and concept integration demonstrate the importance of relatively local analysis of the structure and organization of concepts (Holland, Holyoak, Nisbett, & Thagard, 1986; Gentner, 1989; Fauconnier & Turner, 1998). Processes to identify, select and compare sub-networks of concepts are common to these models of complex patterns of reasoning. The models also support the present view that good conceptual understanding relies on a rich well-structured network of concepts.

Research on expertise in physics and mathematics problem solving also speaks to this issue. Expert knowledge is usually organized in accordance with the important core principles of the domains (Larkin, McDermott, Simon, & Simon, 1980; Chi, Feltovich, & Glaser, 1981; Koedinger & Anderson, 1990). So, approaches to conceptual learning should foster the principled organization of concepts. It is, however, unrealistic to expect most learners to become full experts of domains in science and mathematics, given their scope. Thus, a good understanding should: (1) encompass a fair initial sample of concepts from throughout the whole domain, so that when novel situations are met there are likely to be some relevant concepts to act as an appropriate starting point for problem solving; (2) include problem solving methods that are easy enough for the learners to use to construct the new concepts for themselves.

These considerations provide some indication of the possible goals of conceptual learning in terms of building good networks of concepts. Support for learning should facilitate the development of networks that: have ample scope; are well integrated under some overarching scheme that reflects the fundamental laws of the domain; have good differentiation of concepts as they relate to different classes of situations; provide usable methods for independent analysis of new concepts.

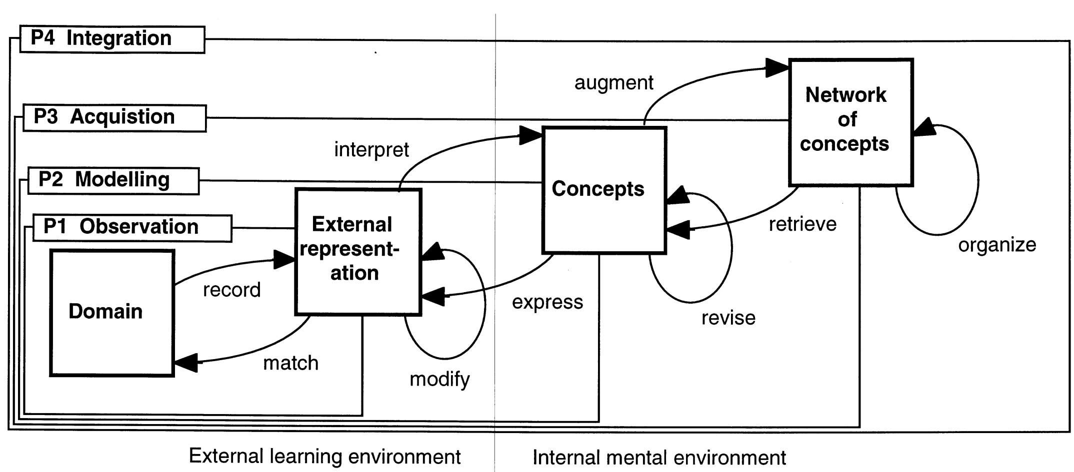

Fig. 3. Components, processes and sub-processes of conceptual learning.

2.2. Components and processes of conceptual learning

How do learners obtain such conceptual networks? Fig. 3 illustrates the major task facing the learner and also the designer of effective approaches to instruction, by showing major gaps to be spanned between the external sources of knowledge exhibited in books, instructors, computer software and as physical models and demonstrations (Domain — bottom left box) and the internal Network of concepts the learner is gradually developing (top right box). Building the conceptual network clearly does not occur by simply transmitting the knowledge of the domain to the learner. At a minimum, consideration must be given (1) to the role of the External representations used for the domain and (2) to the role of individual Concepts (schemas, sets of related propositions or groups of rules), as shown by the middle boxes in Fig. 3. These components must be considered due to the limitations of human information processing. First, external representations are essential given the complexity of domains and the quantity of the information that must be considered. External media act as memory aids and information processing tools. Forms of external representations include charts, graphs, diagrams, equations, tables and all sorts of formal and semi-formal notations. Second, individual concepts are considered on the assumption that learning is an incremental process in which Networks of concepts are gradually constructed by assembling rules, schemas or chunks that are often considered in relative isolation from each other initially.

Four main processes operating over the components during learning and problem solving can be identified. The processes are shown in Fig. 3 by the nested rectangles (P1–P4) surrounding particular components (domain, external representation, etc.). The arrows between squares for the components represent sub-processes.

(P1) *Observation* generates and checks expressions that are descriptions of phenomena. Its sub-processes include: *record* a phenomenon or a particular case as an expression in the external representations; and, *match* an expression with a phenomenon or case (e.g. drawing a circuit diagram for a new circuit and confirming that the diagram matches the circuit topology).

(P2) *Modelling* subsumes observation and acts to generate new expressions to tie together descriptions of particular cases or phenomena with selected relations or concepts. It has sub-processes to *modify* an expression to obtain a new expression and to *express* a concept in the external representation. An example is finding a formula and calculating the overall resistance using the diagram for the new circuit (drawn above) by combining known formulas for the resistance of simpler circuits.

(P3) *Acquisition* involves constructing a new concept (mentally). It subsumes modelling and combines it with sub-processes to either *interpret* an expression as a concept, or to *revise* an expression for a concept, or both, in the context of related concepts *retrieved* from the network. An example is forming a chunk of knowledge consisting of the new formula for resistance (see above) and the diagram for a given circuit, as distinct from formulas for other previously analyzed circuits.

(P4) *Integration* is a process that adds a concept to the network of concepts or that modifies the structure of an existing network of concepts in some way. It subsumes acquisition and combines it with sub-processes that *augment* a network of concepts with a new concept or to (*re*)*organize* a network by connecting, moving or deleting concepts. For example, examining how the new concept/chunk (obtained above) is similar to, or differs from, other concepts for closely and distantly related circuits.

There is a hierarchy of processes in this model. To better integrate or enrich the network of concepts by differentiating different concepts (P4), it is necessary to recall, compare and contrast particular concepts. This may require the acquisition (P3) of new information about the chosen concepts, which will normally involve modelling (P2) the general situation to which the concepts apply, so that the structure of expressions of the concepts can be examined. The modelling may in turn need specific data to be found by observation (P1) of the phenomena or by accessing available information about the domain to help constrain the derivation of appropriate expressions. Clearly, conceptual learning is complex. It encompasses nested groups of processes working on the components at different levels and over different time scales. Later in this special edition, Gobet and Wood consider the processes of acquisition and integration in more detail.

PB's solution to the particle collision problem, shown in Fig. 1, illustrates many of these processes and provides an example of a failed attempt at conceptual learning. If PB had been successful he would have augmented his network of concepts with a new instance of a collision that was consistent with the conservation laws of the domain. PB correctly retrieved the conservation laws from memory as part of the process of trying to acquire a new concept but he failed to find an expression that could be usefully interpreted. This was, in turn, due to the difficulties he had with the modelling process. Although the conservation laws were correctly expressed, Steps 2 and 4, PB was unable to do the necessary modifications to the expressions. This was in part due to his failure to examine all the alternative solutions and also because he did not attempt to match this answer to the given problem conditions. This is surprising, as he had initially recorded the details of the problem in the external representation, an example of the observation process.

The ease with which a coherent network of concepts develops will depend on how difficult each of these four main processes are to perform. When the processes do not facilitate transitions between components of the model they can be viewed as potential barriers that have to be overcome by some means. An effective general approach to promoting conceptual learning will provide support for each of the processes by minimizing the barriers or by providing the learner with the means to bypass them.

2.3. Supporting conceptual learning

Various approaches to supporting conceptual learning can be considered as attempts to overcome the barriers to learning identified in the descriptive model presented in Fig. 3. Although computer support for such learning is now common, there appears to have been no dramatic improvement in conceptual learning, because programs typically support just a few of the processes in Fig. 3.

For example, simulation environments, such as Interactive Physics (Knowledge Revolution, 1996) or Electronics Workbench (Interactive Image Technologies, 1993), provide good support for the process of observation, allowing the user to easily examine a wide variety of cases.

Virtual electronic circuits are easily built in Electronics Workbench and their behaviour accurately simulated under different operating conditions. The system provides meters and graphs for the automatic presentation of data. This in turn facilitates the process of modelling to an extent, but leaves the problem of generating appropriate expressions consistent with the relevant laws as an unsupported task for the user. Such systems have extensive simulation engines based on the laws of the domain but these are typically hidden from the learner. There is normally no provision in the systems to demonstrate how the laws constrain particular phenomena or shape the overall structure of the domain.

Providing learners with the laws as mathematical formulas for modelling the domain leaves them with considerable work to express the target concepts in some interpretable form in the external representation. Some of this effort can be mitigated by the use of spreadsheets and symbolic "algebra calculators", but this still leaves unresolved the problem of the acquisition of appropriate concepts. Being able to easily manipulate the formulas algorithmically does not necessarily make the interpretation of the resulting expressions meaningful or any easier to do. Learners can easily lose sight of what complex expressions mean and become narrowly focused on processing formal notation as an end in itself.

An approach that encompasses more of the processes in Fig. 3 is White's (1989, 1993) use of *intermediate causal models*. These are computer-based visualizations at an 'intermediate' level of abstraction between concrete descriptions of phenomena and abstract general laws. In terms of the model in Fig. 3, intermediate causal models have been designed to support the observation, modelling and acquisition processes (P1–P3) by presenting visualizations that show critical information about the phenomena in a form that reflect aspects of particular concepts. However, the approach does not directly support the process of concept integration (P4). The visualizations provide information about particular properties and relations, but it is largely left to the learner to forge connections between them.

Kaput (1992) sees computers as providing an opportunity to exploit the role of external representations to support learning better. The possibilities include, for example; off-loading students' mental computation onto the computer freeing them to focus on meaningful aspects of problems; using the system to represent processes as separate symbols, so that procedures themselves can become objects of reflection. In terms of the model in Fig. 3, these approaches are attempting to enrich the modelling process (P2).

Some experimental systems that aim to promote scientific reasoning skills may support the process of conceptual integration (P4); for example, Belvedere by Suthers and colleagues (Cavalli-Sforza, Moore, & Suthers, 1993) and Convince Me by Ranney and colleagues (Ranney & Schank, 1998). The latter system exploits Thagard's (1989) model of explanatory coherence, mentioned above, to assess the acceptability of the structures of propositions that learners propose as explanations of particular problems. By stating concepts in natural language, rather than the representations specific to particular domains, the systems allow learners to closely examine how concepts interrelate, without needing to invoke the acquisition (P3) and modelling (P2) processes.

In terms of our model of conceptual learning, previous approaches to computer support for learning appear rather piecemeal. None attempt to cover the full range of processes identified in Fig. 3, so the potential benefits of an approach that facilitates one or two processes may be hindered by the lack of support for the other processes. What appears to be required is an approach in which all of the processes involved can be equally supported in a general systemic fashion. This is what the representational analysis and design project is attempting to do.

2.4. Role of representations in conceptual learning

The history of science is rich in examples where a new representation had to be adopted or invented in order for new discoveries to be made (e.g. Miller, 1986; Gooding, 1996). By analogy, if the representations that a scientist uses to make discoveries are fundamental, then the representations given to learners finding out about the same domain for themselves are also likely to be critical. Further, it is possible that the representations used by the scientists may be better suited for students learning the same domain than the representations currently employed, because the original representations are sometimes well tuned to the cognitive requirements of reasoning and problem solving in those domains. However, algebra is now conventionally used, but it is too powerful and unconstrained a representation to be used for straightforward learning (Cheng, 1996c).

Work in cognitive science has clearly demonstrated the importance of understanding representations to understand problem solving (Newell & Simon, 1972). A poor representation may lead to an order of magnitude increase in the difficulty of finding a solution to a problem compared to a good representation (Kotovsky, Hayes, & Simon, 1985). The research in this area may be classified using the model in Fig. 3. There has been work that has addressed the modelling process (P2), in particular examining the nature of external representations. Diagrammatic representations can offer substantial benefits, because of the way that they index information spatially (Larkin & Simon, 1987). Multiple external representations may facilitate or hinder learning, depending on how information is distributed among the representations (e.g. Ainsworth, Bibby, & Wood, 1997; see also the paper by Ainsworth in this volume). The nature of the acquisition process (P3) and the form of concepts has been examined. Visual representations in the form of perceptual chunks and diagrammatic configuration schemas have a fundamental role in expert knowledge of some domains (Koedinger & Anderson, 1990), including electricity (Egan & Schwartz, 1979). Further, the importance of considering both external and internal aspects of representations in terms of how information is distributed in different ways between mind and environment is crucial in this area (e.g. Zhang & Norman, 1994). Logical inferences that require great deliberation may be "off-loaded" on to an external representation, where visual processes can be exploited with much less effort (e.g. Scaife & Rogers, 1996; Kaput, 1992).

Such previous approaches to supporting conceptual learning can be interpreted as attempts to mitigate the problems of the representations employed. Observation (P1) will be harder to do with representations that do not provide clear and simple mappings between the expressions of the model and the elements of the domain. This may in turn require that additional representations be employed, such as charts and graphs. Similarly, modelling (P2) will be more difficult if the expressions in the representation cannot be simply interpreted in terms of laws governing the behaviour of phenomena. For instance, when arbitrary syntactic rules are used to encode the laws of the domain, the expressions in the external representation are unlikely to directly reflect the structure of phenomena or organization of the domain. Acquisition (P3) will be more difficult with a representation that does not effectively allow different concepts to be distinguished easily from each other. Remembering a new concept will be hard if it is not well differentiated from similar but distinct concepts, or if the expression for the concept is similar to expressions for known concepts that are quite unrelated. Conceptual integration (P4) is made more difficult by any representation that does not allow the relations between concepts to be easily explored, or which fails to provide a clear overarching conceptual scheme to support the interpretation of concepts. In other words, a good representation should allow the learner to quickly obtain a sense of the overall topology of the network of concepts and some means to locate where in the structure any particular concept belongs.

Thus, the role of representations in learning is fundamental and it should be possible to substantially improve learning by replacing poor representations with ones that match the criteria just outlined. The question is, now, what constitutes an effective representation for learning?

3. Nature of effective representations

Ideally, an effective representation will not hinder the learner by presenting barriers between any of the pairs of components in Fig. 3. The representational analysis and design project is studying the characteristics of effective representations that support all of the processes in a systemic fashion, as a way to unlock conceptual learning. The particular class of representations that has been the focus of the project are LEDs, because they seem to be effective for conceptual learning. The project has taken various approaches to the study of LEDs, including: the analysis of the role of LEDs in the history of science (Cheng, 1996a); the design of new LEDs (Cheng, 1999a–d); the empirical evaluation of LEDs in comparison to conventional representations (Cheng, 1996c,d, 1999e); the construction of computer-based learning environments that exploit LEDs (Cheng, 1996b, 1998a, 1999b); the theoretical and empirical examination of the internal knowledge structures acquired when learning with LEDs (Cheng, 1998b, 1999a). From these converging lines of research, characteristics of effective representations, which support all the processes of conceptual learning in a systemic fashion, have been derived (Cheng, 1999e, 1999f).

In this section the characteristics are presented, and an example of problem solving with LEDs for electricity is examined to show why the possession of the proposed characteristics may allow representations to effectively promote conceptual learning.

3.1. Characteristics of effective external representations

Five characteristics of representations have been found and they are grouped depending whether they: (1) concern the *semantic transparency* of representations; that is, the ease of interpreting the expressions of the representations in terms of the domain (characteristics C1 to C3); or (2) deal with the usability of representations, *plastic generativity*, that is the ease of forming and manipulating expressions in the representations (characteristics C4 and C5).

(C1) The first characteristic of the semantic transparency group proposes that an effective representation must help to integrate information about a domain at different levels of abstraction. Expressions in a representation should reveal the nature of the connection between the laws of the domain and the descriptions of phenomena.

(C2) The second characteristic in this group proposes that a good representation will simultaneously differentiate concepts that are different in detail, but will show how at a more general level the concepts are related. On the one hand, the invariants of a domain, such as underlying laws or fundamental axioms, should be apparent as universal structural features that are constant across all the expressions of the representation. On the other hand, things which are contingently variable in a domain, such as alternative cases or values of properties, should be reflected as features of the representation which themselves vary across different expressions. The representation of concepts will be globally homogeneous but locally heterogeneous, with respect to the domain.

(C3) The third characteristic is to support the integration of the many different perspectives or ontologies applicable to the domain within a representation. Alternative perspectives should be supported by different readings or interpretations of the same expressions within the representation. New expressions should not have to be generated, nor supplementary representations introduced, to allow complementary views of the same relations or phenomena to be examined.

(C4) The first characteristic in the plastic generativity group suggests that the expressions of an effective representation will be malleable, that is not too brittle nor too fluid. A brittle representation is one in which it is not possible to easily derive all the acceptable expressions that are meaningful in terms of the laws or phenomena of the domain. When a representation breaks down in this sense, some other representation will need to be invoked. A fluid representation is one in which it possible to generate many correct but largely meaningless expressions, as is the case with algebra. To use a fluid representation requires additional constraints to be invoked to make modelling practical; for instance, prior knowledge about the suitability of certain types of derivations for particular classes of problems. Malleable representations are likely to be the most fruitful in searches for appropriate expressions.

(C5) The second characteristic in this group proposes that the procedures in the representation for working with expressions should be compact and uniform. The same procedures for manipulating the representation should be applicable to all the situations and problem types of the domain. Each procedure should require few inferential steps to complete. These attributes will reduce the likelihood of making errors during representation manipulation and also reduce learning demands by requiring mastery of only a few short procedures.

The utility and validity of the characteristics are being demonstrated by using them to design novel representations for domains such as probability theory (Cheng, 1999f). In the next section, an example of a representation that possesses these characteristics is presented and the reasons why possession of the these characteristics makes a representation better for conceptual learning are considered.

These characteristics bear some resemblance to Green's cognitive dimensions of notations (Green & Petre, 1996). The two approaches differ in that Green is concerned with the general usability of representations, whereas the representational analysis and design project is concerned with what makes effective representations for learning.

3.2. Law encoding diagrams

A system of LEDs for a particular domain uses diagrammatic constraints to encode the laws of the domain in the structure of diagrams in such a way that each instantiation (drawing) of a LED represents one instance of the phenomenon and one case of the laws. Examples of systems of LEDs include Galileo's diagrams for kinematics, Newton's diagrams for dynamics, some of the diagrams for chemical bonding and electronic structure, and also Huygens's diagrams for particle collisions (Cheng, 1996b). Here, we will consider LEDs for basic electricity, to illustrate the nature of this general class of representational systems and to show how representations can meet each of the desirable characteristics of effective representations.

The LEDs for electricity are called AVOW diagrams after the basic electrical properties (Amps, Volts, Ohms, Watts) and this representational system was developed as part of the present project (Cheng, 1999d). Fig. 4 shows a simple electrical circuit composed of a number of resistors connected to a battery (left) and the AVOW diagram matching the circuit (right). Each rectangle, or AVOW box, represents one resistor in the network. The height of the box is the voltage (V) across its resistor and the width is the current (I) through the resistor. The gradient/slope of the diagonal across the box gives the resistance (r) of the resistor. The area of the box is the power (P) dissipated. The geometry of rectangles is being used to encode the basic laws of electricity (i.e. $V=Ir$, $P=VI$). Resistor-A has about double the voltage of resistor-B but about the same current. The spatial arrangement of boxes in the AVOW diagram reflects the topology or configuration of the circuit. Resistor-B and resistor-C are in series with each other and together they are in parallel with resistor-A, and all three are as a whole in series with resistor-D. The electrical properties of the network, or any sub-network, are represented by the same diagrammatic properties of the rectangle for the whole network.

Empirical evaluations of learning with AVOW diagrams (Cheng, 1999e) and other LEDs (Cheng 1996b,d) have shown that this class of representations can promote conceptual learning better than conventional representations for the same domains. For example, in the first experiment on electricity a mini-curriculum was developed, which students followed using AVOW diagrams or a conventional algebraic approach (Cheng, 1999e). After about 100 min of instruction the AVOW diagram group had obtained a significantly better understanding of the

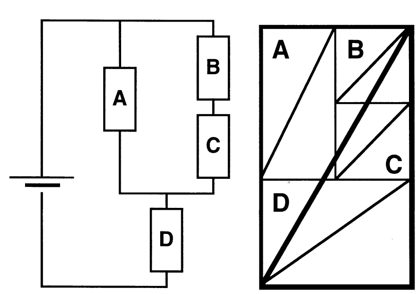

Fig. 4. A circuit diagram and an AVOW diagram.

domain and were able to solve difficult problems which they had not previously seen. The algebra group improved in places but not on the more conceptually demanding problems. Detailed analysis of the learners' problem solutions provided evidence of the types of processes that appear to be responsible for their learning, or their lack of it (Cheng, 1999e). It appears that learners were using the AVOW diagrams to learn concepts as perceptual chunks based on the diagrams which they had been given or which they produced themselves. The perceptual chunks seem to be organized in the form of a lattice, with concepts lower down being based on diagrams composed of components from concepts learnt earlier and positioned higher in the lattice of concepts. The algebra group did not appear to have a well organized set of concepts, but a loose collection of various rules, formulas and propositions.

Fig. 5 gives a problem that was used in that empirical evaluation. This is a deceptively difficult problem to solve accurately, because it involves a complex interaction between the properties of the components in the network. The problem solver must initially work out how the circuit functions correctly, then how it works after the fault develops, and finally compare the details of the two states. We will contrast the post-test solutions of the learners trained using AVOW diagrams or algebra.

Fig. 6 (top) shows the few equations written by the one and only participant in the algebra group who successfully solved the problem. The transcript of the explanation he provided with the answer, shown in Fig. 6 (bottom), has been annotated to show when the participant was considering the basic electrical properties. Each property is considered in relative isolation from the others during the course of the explanation (the *P, V* and *I* symbols tend to be clustered together). Clearly the participant is having difficulty dealing with all the relations and interactions which have to be simultaneously satisfied. Although he got the correct answer, there are errors in his reasoning (as indicated by the underlined section towards the end of the transcript). He had difficulty modelling the situation (process P2, Fig. 3). Although he wrote appropriate expressions for the laws governing the domain, their application to the problem was indirect and clumsy.

This participant has worked out the initial and final states of the situation, so it is possible that he will have acquired a concept encapsulating the behaviours of this network under related conditions (P3). However, the new concept is likely to be fragile as the reasoning behind it is not strictly correct. The error is likely to cause confusion in the future, when similar circuits are again considered. The concept is also going to be poorly integrated with other concepts (P4), because the problem situation was described as a fault in the network, so the broken bulb was simply ignored rather than considered as an integral component of the

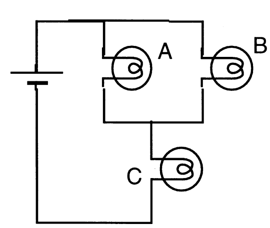

In this circuit what happens to the relative brightness of the light bulbs before and after bulb-A burns out?

Fig. 5. Faulty circuit problem.

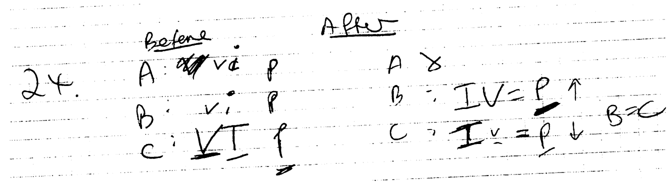

. . . I think B brighter, Cdimmer.

That's because before you made the change to the circuit you've got C receiving all the current **(I)** in the circuit flowing through it and I'd say slightly more than half the voltage **(V)** the source is supplying. Because again the equivalent circuit... the resistance **(r)** of the whole of the parallel thing put together as one in series with C would be slightly less than C itself from the total resistance **(r)** for parallel resistors equation. So therefore C'd draw slightly more, sorry, there'd be slightly more voltage **(V)** across C originally than there would be across each **(V)** of the branches in the parallel circuit where A and B are. So previously before the switch A and B'd have fairly low powers **(P)** but equal **(P)** to each other and C'd have a fairly big power **(P)**, so C'd be brighter **(P)** than those two which would be the same as each other but dimmer **(P)**. Whereas after the change you'd break the circuit at A so A is not shining **(P)** at all now. That would mean that you now set up a situation where B and C are just simply wired in series together to the pack... the power pack... so they'll have equal currents **(I)**. Which means that previously, that would mean that B, the current **(I)** received by B is doubled because previously the whole current **(I)** in the circuit had to be shared between the two identical bulbs, so each has got a half **(I)**, whereas now B's getting all of the current **(I)** as is C's still doing the same **(I)**. And again, because they're in series and both the bulbs are identical they'll have the same resistances **(r)** so they'll each draw the same voltage **(V)**, so they'll each have half the volts **(V)** from the source through them now. So overall because C had a slightly more higher resistance **(r)** than the equivalent of the whole parallel section together, it was previously having a bit more volts **(V)** than the two branches of the parallel, it'll be receiving the same **(V)** now, so C's got slightly less voltage **(V)** than before. Whereas B'll have a bit more now **(V)** because it had slightly less **(V)** than half before and it's got exactly half **(V)** of it now.

So overall, with B receiving more current **(I)** and more voltage **(V)** it's power'd **(P)** go up so it'll get brighter, and with C now receiving the same current **(I)** but slightly less voltage **(V)** it's power **(P)** will go down so C'll get dimmer. That's it.

Fig. 6. An conventional solution of the faulty circuit problem.

circuit whose electrical properties happened to have changed. There is no indication in the transcript that the participant sees the connection between the broken bulb and the closely related phenomenon of insulators.

Fig. 7 is an AVOW group participant's solution to the same problem (it has been redrawn for clarity). Initially this participant draws the AVOW diagram for the circuit functioning correctly [Fig. 7 (before)]. The diagram shows the correct configuration of AVOW boxes for the given circuit topology and also correctly encodes the identical resistances of the bulbs as rectangles with equal gradients. The areas of the boxes shows that bulb-A and bulb-B shine

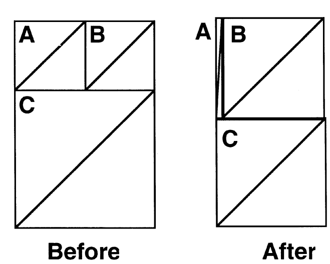

Fig. 7. AVOW diagram solution to faulty circuit problem.

equally brightly but are substantially dimmer than bulb-C. In AVOW diagrams an insulator is represented by a vertical line segment, so in Fig. 7 (after) the AVOW box for bulb-A becomes a becomes a tall thin box. The overall shape of the box for bulb-B is kept the same but its size is increased now that bulb-A has blown. However, the participant realizes that the overall height of the diagram must be constant and that the shape of the box for bulb-C must also be kept constant. Thus, the size of the box for bulb-C is reduced, so that all of these constraints can be met. The participant then simply reads off the correct answer by comparing the relative sizes of the boxes in the before and after diagrams in Fig. 7 (bulb-B gets brighter and bulb-C gets dimmer than before, but after they are equally bright).

Modelling (P2) is relatively straightforward for the AVOW participant. The diagrams were drawn straight off for the two states. The participant has clearly acquired (P3) perceptual chunks for series and parallel sub-networks, which are used for modelling. A new chunk for each diagram can be easily acquired for each diagram and taken together they may be integrated (C4) into a large chunk for a typical fault situation. Note, also how the concept of a broken bulb is integrated within the representation with concepts about insulators or ‘‘open’’ circuits.

3.3. Conceptual learning with LEDs

The electricity example has shown how LEDs can support the main processes of the model in Fig. 3. Now, the observations and findings of the studies of learning with AVOW diagrams and other LEDs will be interpreted in terms of the characteristics of effective representations and the model of conceptual learning in Fig. 3.

AVOW diagrams, like LEDs in general, integrate levels of abstraction (C1) by incorporating into the same diagrams (1) information on the magnitudes of properties of specific components represented as the size of the diagrammatic features and (2) relations governing the domain in the form of diagrammatic constraints. In the AVOW diagram solution to the faulty circuit problem (Fig. 7), the learner has used this feature of LEDs to good effect. The diagrams not only have the right arrangement of boxes but are also drawn to the scale, so that direct comparison between magnitudes of particular properties can be made. Like White's (1989) intermediate causal models, LEDs support the processes of modelling (P2) and acquisition (P3). But they do so by reducing the cognitive distance between descriptions of phenomena and abstract laws by combining the levels of information in a single representation, characteristic C1, rather than introducing an additional layer of representations to mediate between the levels.

LEDs typically support the globally homogeneous and locally heterogeneous representation of concepts (C2). At a global level, AVOW diagrams provide a coherent space within which all the components of electrical circuits can be distinguished. The space uses the vertical dimension for voltage and the horizontal dimension for current. This has benefits for learning in terms of the observation process (P1), because the general scheme provides constraints on how to record information in expressions/diagrams of the external representation. At a higher level, this characteristic means that AVOW diagrams support the process of acquisition (P3) of concepts. For example, on seeing the arrangement of AVOW boxes for parallel loads it is easy to spot that the form of the diagram is different to the configuration for series loads, and that quite different relations amongst the voltages and currents must hold for the two cases. Similarly, this characteristic (C2) supports the process of conceptual integration (P4). For instance, the perceptual chunks (concepts) for batteries and loads are similar, as both are boxes. However, they differ in detail with the direction of flow of the current; up the diagram for the battery, because it supplies voltage, and down the diagram for a load. Thus, a new concept can be obtained by combining these two perceptual chunks: a complete circuit is conceptualized as a loop, composed of an AVOW box for the battery connected to the top and bottom edges of the diagram for a network of resistors.

LEDs permit multiple perspectives to be integrated (C3). The use of geometric relations in AVOW diagrams to encode the laws of electricity is consistent with algebraic statements of the laws, but AVOW diagrams may also be interpreted in concrete physical terms. An AVOW box for a resistor may be viewed as a stream of current running downwards and contained by the sides of the box. In the limit, a vertical line down the AVOW diagram may be considered as the path followed by a single electron. The height of the box represents the force driving the flow. The integration of different perspectives (C3) is achieved in LEDs by permitting multiple interpretations of the same diagrams. One of the benefits of this is the support of the process of acquisition (P3), because the different perspectives will provide mutual constraints on each other. The triangulation of alternative interpretations around the single representation will help to clarify and elaborate concepts under each perspective.

AVOW diagrams are relatively easy to use for observation (P1) and modelling (P2), because: (1) they have an appropriate degree of malleability (C4) in the form of a sufficient, but otherwise limited, set of operations for generating correct diagrams; (2) the procedures for manipulating the diagrams are compact and uniform (C5). This is clearly illustrated by comparing Figs. 6 and 7. The diagrammatic solution is typical of the learners in the AVOW group. They have little trouble choosing and applying correct geometric rules for the construction of the diagrams; normally they draw the correct form straight off or make a few incremental modifications to a diagram. There is typically one symbol for each element in the domain, a line for an electrical property and a box for a circuit component. Thus most of the work on the diagrams can be directed towards achieving problem goals, rather than being used simply to marshal related symbols into canonical forms suitable for interpretation. The algebraic solution was generated by the only learner in the other group to give the correct answer. In other respects this solution is typical, in that it shows a reluctance to attempt to model the circuit by writing formulas and computing values. Learners prefer to, or have to, resort to informal verbal reasoning, as in this case, which happens to show the consequence of using a representation that is too unconstrained, flexible, and certainly not compact, Fig. 6.

In summary, the AVOW diagram representation possesses each of the desirable characteristics of effective representations. This in turn means that the properties of the representation seem to facilitate all of the processes involved in conceptual learning identified in Fig. 3. Thus, by changing the representation, it is possible to reduce many of the barriers to learning under a single systemic approach. This systemic approach is in marked contrast to the previous approaches, discussed above, that mostly attempt to enhance conceptual learning by addressing one or two of the processes in Fig. 3, rather than the process as a whole.

4. Education with new representations

This paper has summarized some of the work of the representational analysis and design project. The central theoretical claims are that: the representations used for learning can substantially affect what is learnt and how easy learning occurs; representations can constrain the nature of the conceptual structures that the learners develop and the problem solving procedures they acquire. Thus, serious consideration must be given to the ontological, structural and functional roles that representations have in learning. By doing this, the project is demonstrating that alternative representations can unlock conceptual learning.

There are general implications of this work not only for the design of computer-based systems to support learning, but also all forms of learning that require representational systems of any complexity. This final section considers some of the implications of using effective representations like LEDs.

The model of conceptual learning introduced above highlights the complexity, variety and number of processes involved in such learning. It is unlikely that there will be any straightforward solution to effectively promoting such learning. This is why the present project has adopted a systemic approach to promote conceptual learning through the analysis and design of representations designed to provide support for each of the constituent processes involved.

The project has identified characteristics of effective representational systems for learning. Two general consequences follow from the formulation of such characteristics. First, it is feasible to analyse existing representations to predict the nature of the misconceptions about the domain and difficulties in their intended use caused by the representation. The examples of learners using algebra illustrate the point. The second consequence is the feasibility of designing new representations that are less likely to foster misconceptions and that will facilitate problem solving.

Although effective representations such as LEDs hold some promise to enhance learning in particular domains, there are various wider issues that may limit the value of such alternative representations, two of which are considered here.

First, it is desirable for students to become familiar with problem solving using algebra, given the extent to which the approach is embedded into education. Thus, will using LEDs initially to learn about a domain assist or hinder the later transition to solving problems using the conventional representations for that domain? Compared to the other approaches, mentioned above, that advocated the use of informal propositions or realistic visualizations, LEDs may better support the transition into algebra. LEDs encode the laws of the domain using geometric rules which can often be straightforwardly mapped into algebraic formulas. Few of the other approaches attempt to provide support for bridging the conceptual gulf between concrete descriptions and the algebraic equations stating the laws. However, it remains an open empirical question whether well-designed instruction based primarily on algebra will be better than starting with LEDs and then later making the transition.

The second issue that may limit the potential of adopting new and effective representations is their novelty. Any attempt to introduce new representations into curricula in which traditional representations are deeply and comprehensively embedded will be a substantial task. For some educators, understanding a domain is synonymous with knowing the appropriate expressions in the traditional representation, so alternative representations are seen as largely irrelevant or even as undesirable distractions. Other educators may appreciate the potential benefits of alternative representations but not adopt them because the traditional representations are just too tightly interwoven with course materials, forms of assessment, and teachers' knowledge and skills. Alternative representations may require too great a departure from existing resources and practices to seem to be worth adopting. Thus, one of the long term objectives of the present project is to demonstrate that switching representations can improve students' conceptual learning to such an extent that the major changes needed for their introduction will be seen as worthwhile.

Acknowledgements

This work was supported by the UK Economic and Social Research Council through the Centre for Research in Development, Instruction and Training. My thanks goes to the members of the Centre who have assisted and encouraged me in the pursuit of this work, particularly David Wood, Nigel Pitt and Lucy Copeland.

References

Ainsworth, S. A., Bibby, P., & Wood, D. J. (1997). Information technology and multiple representations: new opportunities — new problems. *Journal of Information Technology for Teacher Education*, *6*(1).

Anderson, J. R., Corbett, A. T., Koedinger, K. R., & Pelletier, R. (1995). Cognitive tutors: Lessons learned. *The Journal of the Learning Sciences*, *4*(2), 167–207.

Cavalli-Sforza, V., Moore, J. D., & Suthers, D. (1993). Helping students articulate and criticize scientific explanations. In P. Brna, S. Ohlsson, & H. Pain, *Proceedings of AI-ED 93: World conference on artificial intelligence in education*. Charlottesville, VA: Association for the Advancement of Computing in Education.

Cheng, P. C.-H. (1996a). Scientific discovery with law encoding diagrams. *Creativity Research Journal*, *9*(2&3), 145–162.

Cheng, P. C.-H. (1996b). Law encoding diagrams for instructional systems. *Journal of Artificial Intelligence in Education*, *7*(1), 33–74.

Cheng, P. C.-H. (1996c). Problem solving and learning with diagrammatic representations. In D. Peterson, *Forms of Representation* (pp. 47–66). Exeter: Intellect Books.

Cheng, P. C.-H. (1996d). Learning qualitative relations in physics with law encoding diagrams. In G. W. Cotter, *Proceedings of the eighteenth annual conference of the cognitive science society* (pp. 512–517). Hillsdale, NJ: Erlbaum.

Cheng, P. C.-H. (1998a). Some reasons why learning science is hard: Can computer based law encoding diagrams make it easier? In B. P. Gotten, H. M. Half, C. Redialed, & V. Shute, *Intelligent tutoring systems* (pp. 96–105). Berlin: Springer.

Cheng, P. C.-H. (1998b). A framework for scientific reasoning with law encoding diagrams: Analysing protocols to assess its utility. In M. A. Gernsbacher, & S. J. Derry, *Proceedings of the twentieth annual conference of the cognitive science society* (pp. 232–235). Hillsdale, NJ: Erlbaum.

Cheng, P. C.-H. (1999a). Networks of law encoding diagrams for understanding science. *European Journal of Psychology of Education*, *14*(2), 167–184.

Cheng, P. C.-H. (1999b). Interactive law encoding diagrams for learning and instruction. *Learning and Instruction*, *9*(4), 309–326.

Cheng P.C.-H. 1999c. *Visualizing classical particle interactions: composing diagrams to model collisions* (Technical No. 60), ESRC Centre for Research in Development, Instruction & Training.

Cheng P.C.-H. 1999d. *AVOW diagrams: A representational system for modelling electricity* (Technical No. 61), ESRC Centre for Research in Development, Instruction & Training.

Cheng P.C.-H. 1999e. *Electrifying representations for learning: An evaluation of AVOW diagrams for electricity* (Technical No. 62), ESRC Centre for Research in Development, Instruction & Training.

Cheng P.C.-H. 1999f. *Representational analysis and design: What makes an effective representation for learning probability theory?* (Technical No. 63), ESRC Centre for Research in Development, Instruction & Training.

Cheng, P. C.-H., & Simon, H. A. (1995). Scientific discovery and creative reasoning with diagrams. In S. Smith, T. Ward, & R. Finke, *The creative cognition approach* (pp. 205–228). Cambridge, MA: MIT Press.

Chi, M. T. H. (1992). Conceptual change within and across ontological categories: examples from learning and discovery in science. In R. N. Giere, *Cognitive models of science* (pp. 129–186). Minneapolis: University of Minnesota Press.

Chi, M. T. H., Feltovich, P. J., & Glaser, R. (1981). Categorization and representation of physics problems by experts and novices. *Cognitive Science*, *5*, 121–152.

Egan, D. E., & Schwartz, B. J. (1979). Chunking in recall of symbolic drawings. *Memory and Cognition*, *7*(2), 149–158.

Fauconnier, G., & Turner, M. (1998). Conceptual integration networks. *Cognitive Science*, *22*(2), 133–188.

Gentner, D. (1989). The mechanisms of analogical learning. In S. Vosniadou, & A. Ortnoy, *Similarity and analogical reasoning* (pp. 199–241). London: Cambridge University Press.

Gooding, D. C. (1996). Scientific discovery as creative exploration: Faraday's experiments. *Creativity Research Journal*, *9*(2 & 3), 189–206.

Green, T. R. G., & Petre, M. (1996). Usability analysis of visual programming environments: a 'cognitive dimensions' framework. *Journal of Visual Languages and Computing*, *7*, 131–174.

Hiebert, J., & Carpenter, T. P. (1992). Learning and teaching with understanding. In D. A. Grouws, *Handbook of research on mathematics teaching and learning* (pp. 65–97). New York: Macmillan.

Holland, K., Holyoak, K., Nisbett, R., & Thagard, P. (1986). *Induction: Processes of inference, learning, and discovery*. Cambridge, MA: MIT Press.

Interactive Image Technologies 1993. *Electronics workbench user's guide version 5*. Toronto: Interactive Image Technologies.

Kaput, J. J. (1992). Technology and mathematics education. In D. A. Grouws, *Handbook of research on mathematics teaching and learning* (pp. 515–556). New York: Macmillan.

Knowledge Revolution (1996). *Interactive physics users manual version title* > . San Mateo, CA: Knowledge Revolution.

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

Kotovsky, K., Hayes, J. R., & Simon, H. A. (1985). Why are some problems hard? *Cognitive Psychology*, *17*, 248–294.

Kyllonen, P. C., & Shute, V. J. (1989). A taxonomy of learning skills. In P. L. Ackerman, R. J. Sternberg, & R. Glaser, *Learning and individual differences: Advances in theory and research*. New York: Freeman.

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

Larkin, J. H., McDermott, J., Simon, D. P., & Simon, H. A. (1980). Models of competence in solving physics problems. *Cognitive Science*, *4*, 317–345.

Lesgold, A., Lajoie, S., Bunzo, M., & Eggan, G. (1992). Sherlock: A coached practice environment for an electronics troubleshooting job. In J. Larkin, & R. Carboy, *Computer assisted instruction and intelligent tutoring systems: Shared issues and complementary approaches* (pp. 201–238). Hillsdale, NJ: Erlbaum.

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

Newell, A., & Simon, H. A. (1972). *Human problem solving*. Englewood Cliffs, NJ: Prentice–Hall.

Ranney, M., & Schank, P. (1998). Towards an integration of the social and scientific: observing, modelling and promoting the explanatory coherence of reasoning. In S. J. Read, & L. C. Miller, *Connectionist models of social reasoning and behaviour*. Mahwah, N.J: Erlbaum.

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

Skemp, R. R. (1971). *The psychology of learning mathematics*. Harmondsworth: Penguin.

Thagard, P. (1989). Explanatory coherence. *Behavioral and Brain Sciences*, *12*(3), 435–464.

White, B. (1989). The role of intermediate abstractions in understanding science and mathematics. In *Proceedings of the eleventh annual meeting of the cognitive science society* (pp. 972–979). Hillsdale, NJ: Erlbaum.

White, B. (1993). ThinkerTools: Causal models, conceptual change, and science education. *Cognition and Instruction*, *10*(1), 1–100.

Zhang, J., & Norman, D. A. (1994). Representations in distributed cognition tasks. *Cognitive Science*, *18*(1), 87–122.