



The functions of multiple representations

Shaaron Ainsworth

*ESRC Centre for Research in Development, Instruction and Training, School of Psychology, University Park,
University of Nottingham, Nottingham NG7 2RD, UK*

Abstract

Multiple representations and multi-media can support learning in many different ways. In this paper, it is claimed that by identifying the functions that they can serve, many of the conflicting findings arising out of the existing evaluations of multi-representational learning environments can be explained. This will lead to more systematic design principles. To this end, this paper describes a functional taxonomy of MERs. This taxonomy is used to ask how translation across representations should be supported to maximise learning outcomes and what information should be gathered from empirical evaluation in order to determine the effectiveness of multi-representational learning environments. © 2000 Elsevier Science Ltd. All rights reserved.

1. Introduction

Multi-media and multi-representational learning environments are ubiquitous and were so even before the advent of modern educational technology. A common justification for using more than one representation is that this is more likely to capture a learner's interest and, in so doing, play an important role in promoting conditions for effective learning. The main aim of this paper, however, is to consider the different ways in which multiple external representations (MERs) are used to support cognitive processes in learning and problem solving with computers, and to examine critically the view that the use of MERs confers not only motivational benefit but also leads the learner to a deeper understanding of the subject being taught.

The many multi-representational computer-based learning environments used in classrooms

E-mail address: sea@psychology.nottingham.ac.uk (S. Ainsworth).

0360-1315/00/\$ - see front matter © 2000 Elsevier Science Ltd. All rights reserved.

PII: S0360-1315(99)00029-9

today support a variety of learning activities. Design ranges from general tools including spreadsheets and graphing packages to dedicated software such as FunctionProbe (Confrey, 1992), which uses graphs, equations and tables to teach understanding of functions, and the Blocks World (Thompson, 1992) that helps children understand arithmetic by using Dienes blocks as well as the more conventional symbolic representations. The use of such environments seems destined to become even more widespread, but research designed to evaluate how effectively such multi-representational environments support learning has produced mixed results and implies a degree of caution in their use. Although a number of studies have shown that learners find working with MERs very difficult, and have failed to find the promised learning benefits (e.g. Tabachneck, Leonardo & Simon, 1994; Yerushalmy, 1991), others have show that advantages can accrue from their use (e.g. Ainsworth, Wood & O'Malley, 1998b; Cox & Brna, 1995; Thompson, 1992). The *reasons* for such conflicting findings are a focus of this evaluative review.

Despite conflicting findings about the impact of MERs on learning outcomes, one result found consistently across studies is that learners find translating¹ between representations difficult. For example, Schoenfeld, Smith and Arcavi (1993) examined one student's understanding of mathematical functions using the Grapher environment, which exploits both algebraic and graphical representations to support learning. Observed over a number of sessions, the researchers showed how the student's increasingly successful performance suggested that she had benefited from and mastered fundamental components of the learning domain exploiting both algebra and graphs. However, more detailed and critical analyses of performance revealed that the learner had failed to grasp important connections between the two modes of representation. Similarly, Yerushalmy (1991) found that even after extensive experience with multi-representational learning experiences designed to teach understanding of functions, only 12% of students gave answers that involved both the numerical and visual representations. Most answers reflected the use of one representation and a neglect of the other. Such research suggests that appreciating the links across multiple representations is not automatic.

An investigation by Ainsworth, Wood and Bibby (1996) demonstrates how the achievement of translation between one representation and another varies depending upon the nature of the relations between representations selected. The system evaluated, CENTS (see Fig. 1) is designed to teach computational estimation. It uses MERs to focus on the question of how learners understand the relation between the accuracy of their estimates on the one hand and the objective correctness of the answer itself on the other.

Pairs of representations were used to display the direction and magnitude of estimation accuracy. These consisted of either two pictorial representations, two mathematical representations or one pictorial representation and one mathematical representation combined to give a mixed system. The results showed that although children in all experimental groups learnt how to estimate more accurately, only the children given pairs of pictorial or pairs of

¹ Translation throughout this paper is used to refer to **all** cases when a learner must see the relation between two representations. It is used to refer both to the cases when a learner must comprehend the relation between two representations and also when they must act to reproduce this relation. It is neutral about whether translation occurs through direct mapping between the symbols or whether it is mediated through domain understanding.




Fig. 1. CENTS displaying two pictorial representations.

mathematical representations also improved at judging the accuracy of their own estimates. Those who received mixed pairs did not. Each of the representations used in the mixed condition was also present in either the pictorial and mathematical conditions where it was used successfully. The poorer performance associated with mixed representations is likely to lie in the demands of translating between representations, rather than in the properties of the individual representations per se. The issue of how one assesses the achievement of such translations between representations will be returned to later.

To overcome problems in learning how to translate between representations, many learning environments have been designed to exploit automatic translation or “dyna-linking”. Here a learner acts on one representation and is shown the effects of their actions on another. It is hoped that if a system automatically performs the translation between representations, then the cognitive load placed on learners should be decreased and so free them to learn the relation between representations (e.g. Kaput, 1992; Scaife & Rogers, 1996). Against this position, advocates of a constructivist approach to education might argue that dynamic linking leaves a learner too passive in the process. Such dyna-linking may discourage reflection on the nature of the translations leading to a concomitant failure by the learner to construct the required understanding. At present, such global issues cannot be resolved, and this is likely to remain

the case until we understand more about the conditions under which multi-representational learning environments should be designed to support cross-representation translation.

In this paper, it will be argued that MERs are used for several distinct purposes. A failure to acknowledge this explains many of the conflicting findings arising out of evaluations of multi-representational learning environments. Furthermore, it is proposed that generalised principles for effective learning with MERs will rest upon a careful analysis of their purposes. To this end, a functional taxonomy of MERs will be proposed. To illustrate its generality, this will be exemplified by reference to the learning environments described in other papers in this special issue and to cases from the relevant literature. Having identified the different functions underpinning the use of MERs, two further issues are raised. The first is the question of how translation across representations should be supported to maximise learning outcomes. The second considers what information researchers and teachers might gather from empirical evaluation in order to determine the effectiveness of their multi-representational learning environments.

2. A functional taxonomy of multiple representations

A conceptual analysis of existing multi-representational learning environments suggests there are three main functions that MERs serve in learning situations — to complement, constrain and construct. The first function is to use representations that contain complementary information or support complementary cognitive processes. In the second, one representation is used to constrain possible (mis)interpretations in the use of another. Finally, MERs can be used to encourage learners to construct a deeper understanding of a situation. Each of the




Fig. 2. A functional taxonomy of multiple representations.

three main functions of MERs can be further sub-divided into several subclasses (see Fig. 2). Often a single multi-representational environment may serve several of the functions shown, but, to begin with, each class will be considered separately.

2.1. Using MERs in complementary roles

One reason to exploit MERs in learning environments is to take advantage of representations that have complementary roles, where differences between representations may either be in the information that each contributes, or in the processes that each supports. By combining representations that complement each other in these ways, it is envisaged that learners will benefit from the sum of their advantages.

2.1.1. MERs to support complementary processes

The most familiar rationale for using more than one representation is to benefit from the varying computational processes supported by different representations. There is an extensive literature showing that representations that contain the equivalent information can still support different inferences. One common distinction drawn is that between diagrams and sentential representations. For example, Larkin and Simon (1987) proposed that diagrams exploit perceptual processes by grouping together relevant information and hence make processes such as search and recognition easier. Further research has shown that other common representations differ in their inferential power (e.g. Cox & Brna, 1995; Kaput, 1989; Meyer, Shinar & Leiser, 1997). For example, tables tend to make explicit specific values, emphasise empty cells (so directing attention to unexplored alternatives), support quicker and more accurate readoff and highlight patterns and regularities across cases or sets of values. To take another example, the quantitative relationship that is compactly expressed by the equation $y = x^2 + 5x + 3$ fails to make explicit the variation which is evident in an (informationally) equivalent graph, which reveals trends and interaction more directly than an alphanumeric representation.

One example of such a multi-representational learning environment is ReMIS-CL (see Cheng, 1996; this volume). It provides a total of seven different representations to help students understand the nature of elastic collision. Many of these representations are Law Encoding Diagrams (LEDs). An LED is a representation that correctly encodes the underlying relations of one or more mathematically expressible scientific law(s) by means of geometric, topological or spatial constraints such that each instantiation of a LED represents both a representation of any given instance of phenomena or one case of the laws. The representations include numerical equations, one dimensional property diagrams, mass velocity diagrams, and velocity–velocity graphs. In ReMIS-CL an animation of the collision is always available and the learner/instructor can choose up to four of the representations to view simultaneously. Although the majority of the representations express equivalent information, each makes salient different aspects of the situation. For example, the one dimensional property diagram (LHS of Fig. 3) supports the mapping of each element and attribute in the simulation onto a corresponding element in the diagram. Consequently, the lines that represent the initial and final velocities in the LED can effectively be overlaid on the simulation. In contrast, the velocity–velocity graph (RHS of Fig. 3) is consistent with the rule that velocities are always




Fig. 3. ReMIS-CL — an learning environment to teach elastic collisions.

derivable from the simultaneous satisfaction of two separate relations, and the two intersections between the diagonal and the ellipse indicate that two pairs of values can be found in relation to each and every case of the law.

Three main classes of reasons for exploiting multiple representations with different computational properties are found (a) when different learners exhibit preferences for different ones, (b) when the learner has multiple tasks to perform and (c) when using more than one strategy improves performance.

If a learning environment presents a choice of multiple representations, learners can work with their preferred choice. Where learners have varying degrees of experience and expertise with different representations, an appropriate combination leaves each free to select and exploit that with which they feel most familiar. More contentiously, it is often claimed that representational preferences stem not just from experience but are also influenced by more stable individual differences. Factors such as IQ, spatial reasoning, locus of control, field dependence, verbal ability, vocabulary, gender and age have been cited as candidates (see Winn, 1987). A common (although by no means consistent finding) is that lower achieving learners are more likely than their higher achieving peers to benefit from graphical representations of a task (see Cronbach & Snow, 1977; Snow & Yalow, 1982). Various taxonomies of cognitive style have been advanced, but this remains a controversial issue with findings of marked intra-individual differences as well as the proposed inter-individual ones. Thus, there is not necessarily a simple or face-valid relation between supposed cognitive style, representational preference and task performance (e.g. Roberts, Wood, & Gilmore, 1994).

To function effectively in a domain, a learner is typically required to perform a number of

different tasks. There is rarely, if ever, a single representation that is effective for all tasks; rather particular representations facilitate performance on some but not on others. What has been termed the “match-mismatch” conjecture by Gilmore and Green (1984) proposes that performance is most likely to be facilitated when the structure of information required by the problem matches the form provided by the representational notation. Empirical support for this conjecture has been provided by Bibby and Payne (1993) who gave subjects instructions on how to operate a simple control panel device using either (informationally equivalent) tables, a procedure or diagrams. To learn to operate the device, a number of different tasks needed to be performed. These included detection of faulty components and identifying misaligned switches. There were significant interactions between task and representation, and no single representation proved better overall. Thus, participants given tables and diagrams identified faulty components faster than ones provided with information expressed as specific verbal rules. These participants, however, proved faster given the task of deciding which switches were mispositioned. Even in such a simple domain, we can see that MERs can be beneficial by providing representations that fit a task more effectively.

Representations and problem solving strategies also interact. For example, Tabachneck, Koedinger and Nathan (1994) investigated learners solving algebra word problems. They identified six external representations (including verbal arithmetic, diagrams and written algebra) which were associated with four strategies (algebra, guess-and-test, verbal-math and diagram). No single strategy proved more effective than any other. However, where the learner employed more than one strategy, their performance was significantly more effective than that of problem solvers who used only a single strategy. As each strategy had its inherent limitations, switching between them made problem solving more successful. Cox and Brna (1995) report a similar effect when students were observed solving analytical reasoning problems. They found that students used a variety of representations (e.g. logic, set diagrams, tables, and natural language) although the majority of individuals stayed with just one in solving a problem. In only 17% of cases did participants use more than one representation, and this tended to be associated with better performance. Consequently, where learners are given the opportunity to use MERs, they may be able to compensate for any weaknesses associated with one particular strategy and representation by switching to another.

It can be seen that there may be considerable advantages for learning with complementary processes because, by exploiting combinations of representations, learners are less likely to be limited by the strengths and weaknesses of any single one.

2.1.2. MERs to support complementary information

A second reason to use complementary MERs is to exploit differences in the information that is expressed by each. Multiple representations tend to be used for this purpose either in cases where a single representation would be insufficient to carry all the information about the domain, or in cases where attempting to combine all relevant information into one representation would over-complicate the learner's task. In each case, there are two sub-classes of this category (a) where each representation encodes unique aspects of a domain and presents different information and (b) where there is a degree of redundant information shared by the two as well as information unique to each.

2.1.2.1. Using MERs which express different information. Where there is an excess of complex information to convey using MERs allows designers to create representations that are individually simpler and more usable. For example, ‘MoLE’, Oliver and O’Shea (1996) is a multi-representational learning environment which teaches modal logic. One representation is a node and link description of the relation between different modal worlds (LHS of Fig. 4). A modal world is also depicted as a grid of polygons that illustrate the content of each world individually (RHS in Fig. 4). In this example, each representation expresses different domain information and there is no redundancy. It is possible that one representation could have been created that carried all the information. For example, the node and link representation could have included the content grid within each node. Yet, had it done so, the representation would have quickly become cluttered and difficult to interpret when more than a few worlds were displayed. Oliver (1998) analysed the performance of learners working with the two representations and suggests that dividing the information across the two allowed learners to concentrate on different aspects of the task, making the learning goals more achievable.

2.1.2.2. Using MERs to support new inferences by providing partially redundant representations. Rather than provide representations that have completely different information, an alternative is to use MERs that provide some shared information, where partial redundancy of information supports new interpretations of the represented domain. These use of representations is common when one representation is designed to provide functional information (e.g. a functional diagram of a heating system) and the other physical information (e.g. a map of the true positions of radiators, boilers, etc).

A classic illustration of this class of MERs is the problem of finding the quickest route




Fig. 4. The relation and world descriptions representations in MoLe.

between two London Underground stations. The London Underground map designed by Harry Beck in the 1930s shows routes and connections but does not preserve geographical and topological information. The Underground map renders the task of finding the most direct train journey between two stations relatively easy. However, this does not guarantee that the map-reader will find the shortest and quickest route. A train journey between two stations that requires a number of changes might be reachable on foot in a matter of minutes (e.g. Bank to Mansion House — a total of seven stations and two different lines but only 200 m by foot). One adaptive solution to such problems of journey planning would be to integrate the information provided by a street map, which preserves information about geographical distance, with the information in the Underground map which gives train routes but not true distance.

Again, it is possible that a single representation could provide all the necessary information to support the required inference, but at the cost of raising additional problems of interpretation and transparency. By distributing information over such partially redundant representations, multi-representational learning environments can create less complicated artifacts, but then introduces demands for translation and integration — a dilemma that is considered later.

2.2. Using MERs to constrain interpretation

A second use of multiple representations is to help learners develop a better understanding of a domain by using one representation to constrain their interpretation of a second representation. This can be achieved in two ways: by employing a familiar representation to support the interpretation of a less familiar or more abstract one, or by exploiting inherent properties of one representation to constrain interpretation of a second.

2.2.1. Using MERs so that a familiar representation constrains interpretation of a second unfamiliar representation

One rationale for exploiting a familiar representation is to support the interpretation of a less familiar or more abstract one and to provide support for a learner as they extend, or revise misconceptions in, their understanding of the unfamiliar. For example, microworlds such as DM³ (Hennessy et al., 1995) and SkaterWorld (Pheasey, Ding & O'Malley, 1997) provide a simulation of a skater alongside a velocity-time graph (amongst other representations). Two of the well-documented misconceptions that children entertain when learning Newtonian mechanics are that a horizontal line on a velocity-time graph must represent a stationary object and that any negative gradient must entail negative direction. When, for example, children experience a DM³ simulation which shows a skater moving forward at the same time as a dynamic, linked graph of its motion continues to move, they may be led to question and revise their initial conceptions of the graphical representation.

Here, the primary purpose of the constraining representation is **not** to provide new information but to support a learner's reasoning about the less familiar one. It is the learner's familiarity with the constraining representation, or its ease of interpretation, that is essential to its function.

2.2.2. Using MERs so that the inherent properties of a representation constrain interpretation of a second representation

In contrast to these cases, there are situations where an abstract or unfamiliar representation can be exploited to constrain the interpretation of a second representation by exploiting some inherent property. For example, it is argued that graphical representations are generally more specific than sentential representations (e.g. Stenning & Oberlander, 1995). If someone is provided with a representation in a natural language expression such as ‘the knife is beside the fork’, there is inherent ambiguity about which side of the knife the fork has been placed. This is not possible when representing the same world pictorially, since the fork must be shown as either to the left or to the right of the knife (e.g. Ehrlich & Johnson-Laird, 1982). So, when these two representations are presented together, interpretation of the first (ambiguous) representation may be constrained by the second (specific) representation independently of issues of familiarity or experience. In other words, one representation can act to force an interpretation of another one.

This function of MERs can be seen in the design of multi-representational learning environments. For example, COPPERS (Ainsworth et al., 1998b) teaches children about multiple solutions to coin problems. Two representations are used to describe each of the children’s solution in detail (see Fig. 5). The first one is a place value representation (RHS Fig. 5). This provides the user with a representation of how many of each type of coin they used in such a way as to make explicit the arithmetic operations they performed. In the second, a more unfamiliar tabular representation expresses equivalent information (per single row), but the mathematical operations are implicit in the values in the cells and column headings (LHS, Fig. 5). Inherent properties of the tabular representation, however, can constrain interpretation of the place value representation. Answers to coin problems such as ‘ $5p + 10p + 5p + 10p$ ’ and ‘ $5p + 5p + 10p + 10p$ ’ may appear very different to young children who have yet to develop an

PREVIOUS

1p	2p	5p	10p	20p	50p	£1	TOTAL
			3	1	1		70p
		4	3	1			70p
		14					70p
5	5	3	2	1			70p
2	4	2	1	2			70p

NEXT

ANSWERS

Good

one of the answers to this problem is

$2 \times 5 \text{ pence} = 10\text{p} +$
 $1 \times 10 \text{ pence} = 10\text{p} +$
 $2 \times 20 \text{ pence} = 40\text{p} +$
 $2 \times 1 \text{ pence} = 2\text{p} +$
 $4 \times 2 \text{ pence} = 8\text{p}$

70p

intuitive grasp of commutativity. Yet, were these solutions to be displayed in the table, they would appear equivalent as each row does not express ordering information. Therefore, if children learn how to translate between the representations, the likelihood of their grasping or accepting the equivalence of the two different orderings found in the place value representation is increased.

2.3. Using MERs to construct deeper understanding

It has been claimed that exposure to multiple representations leads to deeper understanding. For example, Kaput (1989) proposes that “the cognitive linking of representations creates a whole that is more than the sum of its parts. . . . It enables us to ‘see’ complex ideas in a new way and apply them more effectively”. In this paper, ‘deeper understanding’ will be considered in terms of using MERs to promote abstraction, to encourage generalisation and to teach the relation between representations.

2.3.1. Using MERs to support abstraction

Abstraction is a notoriously slippery term. This paper will restrict attempts at defining it to a consideration of three alternative views.

One use of the term is equivalent to ‘subtraction’, where the emphasis is on detecting and extracting only a sub set of features from the initial representation. For example, Giunchiglia and Walsh (1992) in defining abstraction refer to ‘throwing away details’.

An alternative conceptualisation emphasises re-ontologisation rather than simply subtraction. For example, when children learn to add and subtract with both Dienes blocks and written numerals, and grasp the common underlying patterns and invariants exhibited by actions on quantity expressed in both representations, they come to a more abstracted sense of number and base ten (Schoenfeld, 1986).

A third sense implicates abstraction as reification. Kaput (1989) considers reflective abstraction as the process of creating mental entities that serve as the basis for new actions, procedures and concepts at a **higher** level of organisation. Similarly, Sfard (1991) describes reified understanding as what results when a mathematical entity perceived as a process at one level is reconceived as an object at a higher level. So an algebraic expression such as $3(x + 5) + 1$ can have multiple readings. Initially, it could be read operationally as a sequence of operations (add 5 to the number, multiply by 3 and then add 1). Later, it could be understood structurally as one case of a function, an abstract(ed) object in its own right.

So how might multiple representations encourage abstraction? It is hoped that by providing learners with a rich source of domain representations they will translate or construct references across these representations. Such knowledge can then be used to expose the underlying structure of the domain represented. For example, Dienes (1973) argues that perceptual variability (the same concepts represented in varying ways) provides learners with the opportunities for building such abstractions. Learners can discover invariant properties of a domain in the face of perceptually salient but conceptually irrelevant differences in the appearance of any specific instance: a form of analysis by synthesis. The design of the domain representation for the QUADRATIC tutor, evaluated by Wood and Wood in this special

edition, was directly motivated by Dienes' analysis of the development of abstract understanding in algebra.

Also motivated to evaluate Dienes' ideas, Resnick and Omanson (1987) taught children to add and subtract using both Dienes blocks and written numerals. During a substantial intervention programme, children were given mapping instructions about the correspondence of (or the translation between) these two representational systems. Instruction focused on teaching children to identify parallels and regularities across the two systems (e.g. trading with Dienes blocks, carrying, and borrowing with base 10 algorithms). The objective was to determine if children came to understand how both systems represented equivalent actions on quantities, helping them to construct a more abstract understanding of the structure of base 10 arithmetic. Although the intervention was not particularly successful in meeting these aims, it does provide a well-motivated example of the use of multiple representations to support abstraction.

Schwartz (1995) provides evidence that multiple representations can generate understanding that is more abstract. In this case, multiple representations are provided by different members of a collaborating pair. With two tasks (the rotary motion of imaginary gears, text from biology tasks where inferences must be made), he showed that the representations that emerge with collaborating peers are more abstract than those created by individuals. One explanation of these results is that the abstracted representation emerged as a consequence of requiring a single representation that could bridge both individuals' representations. Thus, detail that was present in individual representations was not incorporated into the joint representation.

2.3.2. Using MERs to support extension

Extension or generalisation can be considered as a way of extending knowledge that a learner already has to new situations, but without fundamentally changing the nature of that knowledge. In contrast to abstraction, extended knowledge does not require re-organisation at a higher level. In cognitive models such as ACT*, for example, generalisation often occurs through variablisation and involves no changes in conceptual structure (e.g. Anderson, 1983).

When considering representations, extension can refer to two different aspects of a learning situation — extending the domains where a given representation is used or extending the way that domain knowledge is embodied to include other representations. The first case of extension can be seen whenever a representation, taught for one purpose or in one domain, is used to serve another. For example, common representations such as tables and graphs might first be taught in the maths classroom. Subsequently, they can be used for representing information necessary to solve problems in physics, geography, economics, etc. However, this type of extension, although common in learning situations, is outside the scope of the present analysis as it concerns the application of a common representation to multiple fields, rather than the use of multiple representations to support learning in a common domain.

The second type of extension is extending domain knowledge *through its expression* in a variety of representations. For example, learners may know how to interpret a velocity time graph in order to determine whether a body is accelerating. They can subsequently extend their knowledge acceleration to such representations as tables, acceleration-time graphs, tickertape etc. This process counts as representational extension if a learner exploits an understanding of

how one representation expresses a concept to gain some understanding of the way in which a second representation embodies the same knowledge.

2.3.3. Using MERs to teach relations among representations

This function of MERs is only subtly different from the cases that have already been considered. Similar to extension, the pedagogical goal is explicitly to teach learners how to translate between representations. However, in this case teaching does not extend from knowledge of one well-understood representation to a second. Instead, two or more representations are introduced simultaneously and learning to translate between them is more of a bi-directional process. For examples, the SkaterWorld environment (Pheasey et al., 1997) presents users with a number of representations simultaneously. A simulation of a skater that is intended to constrain interpretation of other more abstract and unfamiliar representations is always visible. Other representations include tickertape, force arrows, net force indicator, tables of velocity, distance travelled and time elapsed (Fig. 6). In addition, learners can choose one graph from velocity–time, distance–time or acceleration–time graphs. It can be seen that there are a number of representations present at any one time in SkaterWorld. Children learning




Fig. 6. SkaterWorld showing the simulation screen.

Newtonian mechanics with this system spend considerable amounts of time relating these different representations.

The QUADRATIC Tutor (Wood and Wood this volume) teaches pupils with only limited experience of algebra to develop some understanding of the quadratic function. In particular, (following Dienes) it uses the area of squares to make salient the properties of algebraic expressions (and vice versa). QUADRATIC attempts to help children to grasp equivalences across the geometric and algebraic representations. Learners can come to understand, for example, how $x^2 + 2x + 1 = (x + 1)^2$ refers to equivalent sets of graphical representation of the $x + 1$ square (see Fig. 7).

Teaching progresses by allowing children to construct squares of different sizes by mapping elements of the algebraic expressions onto elements of squares that they construct on screen, culminating in an expansion of the general case. This is then repeated for the $(x+n)^3$ and the $(x-n)^2$ cases. The subtle differences between extending representational knowledge and relating representational knowledge is illustrated by the designers' wish that users of the system should be new to algebra. Thus, QUADRATIC teaches them to relate two relatively unfamiliar representations. However, if learners already had a substantial knowledge of algebra, then QUADRATIC could be used to extend this knowledge to the novel situation of explaining the expansion and contraction of geometric squares and cubes.

The goal of teaching relations or translations between representations can sometimes be an




Fig. 7. A screen from the Quadratic tutor.

end in itself. For example, much emphasis is placed on learning how to construct a graph given an equation (e.g. Dugdale, 1982). But often the goal of teaching how two representations are related is to serve some other end. In particular, it is hoped that teaching how representations are related may encourage abstraction. It may well be the case that supporting extension, or teaching the relation between representations, is an initial stage in using MERs. It is hoped that if and when learners master these processes, their knowledge of the representations can then be exploited to serve other ends.

3. Using MERs to support multiple functions

This taxonomy provides a means of articulating and classifying designer's or users' goals in exploiting MERs. However, it has simplified the problem by assuming that a specific multi-representational environment fits into a single category. In reality, many environments embody multiple goals. The way that the functional taxonomy can be used to illustrate this can be demonstrated in reference to COPPERS. The MERs employed as feedback in COPPERS were described above as being used for constraint — the order irrelevance of a row in the table representation could be used to constrain the order sensitive place value representation. However, these representations serve three further functions. Firstly, the place value representation is used to constrain interpretation of the tabular representation. Table representations are often difficult for younger children (e.g. Underwood & Underwood, 1987). COPPERS uses the place value representation to help them understand the information in the table on the grounds that the symbolic procedures of multiplication and addition that are explicit in the place value representations are those required to interpret the tabular representation. The table representation also provides additional information, in that it provides learners with access to a history of previous answers to a question. Finally, the table representation also supports different processes. For example, it emphasises empty cells and makes patterns and regularities in answers more obvious. Moreover, the multiple solutions requested in COPPERS are themselves multiple representations of answers to the problem. The pedagogical objective was that by requiring children to produce several different decompositions to the problem they would start to develop a deeper or more abstracted understanding of the domain. For example, producing non-canonical decompositions (e.g. that 32 is two tens and 12 units) is thought to be crucial in developing conceptual understanding of place value (Resnick, 1992).

This example shows that even relatively simple systems like COPPERS can use MERs in interesting and sophisticated ways. If conflicting and confusing findings about the impact and value of MERs in learning are to be avoided, then the kind of analysis just illustrated is necessary to identify the potential uses and outcomes underlying the use of MERs — the question then becomes how to evaluate and support such uses of MERs.

4. The role of translation in learning with MERs

There are many different reasons why MERs can be beneficial for learning. It was suggested

that MERs are commonly used for one of three main purposes (i.e. that MERs can provide complementary information and processes, can constrain interpretations and help learners construct a deeper understanding of the domain). For each of these uses, multiple sub-components were identified. Furthermore, MERs used in a single system may fulfil two or more of these purposes either simultaneously or sequentially. Identifying the different functions that multiple representations play is crucial as each makes distinct predictions about how the learning goals should be supported. In each case, it is the role of translation between representations which influences the fit between the design and the learning objective(s).

The first use of MERs is to support complementary processes and information. This design is ideal if one representation would be very complex to interpret when it included all of the necessary information. It is also advantageous when the computational properties of alternative representations support and focus on different aspects of the domain or encourage different strategies. As each representation contributes something separate to the process of learning, one way to make the learning task more tractable is to discourage users of a system from learning to translate between representations. This argument is based on the abundant evidence that translating between representations can be very difficult (Tabachneck et al., 1994; Schoenfeld et al., 1993). Furthermore, if translation is not necessary to achieve the particular function of the learning environments representations, then providing co-present representations may encourage learners to attempt to relate them, so inhibiting the achievement of learning outcomes (Ainsworth et al., 1996). Therefore, to maximise this use of MERs, the learning environment should automatically perform translation between the representations, thus freeing the learner from trying to perform this task. Alternatively, it may be appropriate to present the representations sequentially to discourage attempts at co-ordination.

The second category of use of MERs is to constrain interpretation. For example, this can be seen when a known representation supports the interpretation of an unfamiliar abstract representation (e.g. a simulation environment that presents concrete illustrations alongside more abstract representations such as graphs or equations). In contrast to the first use of MERs, designers need to ensure that learners can grasp the relation between the representations. However, the goal of this function of MERs is not to teach learners to translate between representations as this is often a long and complex process. Instead, it is hoped to exploit learners' understanding of the relation between the representations to some further end. This suggests that a learning environment should make very explicit the relation between representations. This could either be achieved by automatic translation or dynamic linking — as a learner manipulates one representation, another one changes (e.g. *Blocks World*; Thompson, 1992). If neither representation is used for action, then ideally a learning environment should signal the correspondence between representations. For example, *COPPERS* indicates the relation between an entry in the table representation and the parallel component of the place value representation by highlighting the corresponding elements in the two representations. Finally, if learners are required to perform this linking of representations for themselves, whenever possible representations that are more easily co-ordinated should be selected. Previously work has identified factors that determine how easily learners can co-ordinate representations (Ainsworth et al., 1998a). The main conclusion can be summarised as “the more that the format and operators of two representations differ, the harder it will be for

a learner to appreciate the relations between them”. Factors which determine the similarity of representations include labelling, modality, interfaces (e.g. direct manipulation or keyboard), strategies, interpretation versus construction of representations, and differences in levels of abstraction.

The third category of MERs is when learners are encouraged to construct a deeper understanding of a domain. It was suggested that this could occur through abstraction, extension or by directly teaching the relation between representations. This goal provides designers with hard choices. If users fail to translate across representations, then deeper understanding may not occur. Yet, representations which provide the different viewpoints normally needed for deeper understanding are those that previously have been shown to be the most complex to relate. Furthermore, educational practice that emphasises the role of the learner in actively constructing their own understanding would suggest that dynamic linking leaves a learner too passive in this process. This over-automation may not encourage users to reflect actively upon the nature of the connection and could in turn lead learners to fail to construct the required deep understanding. The question that remains is what is the best way to achieve the cognitive linking of representations in the mind of the learner.

One approach to this problem is by scaffolding and, in particular, contingency theory (Wood, Bruner & Ross, 1976; Wood and Wood, this volume). This approach suggests that the level of support provided to the learner for any given task should vary depending upon their performance. As a learner succeeds, support should be faded out, but upon failure, then the learner should receive help immediately. In order for a learner to achieve the cognitive linking of representations, the strategy suggested by scaffolding is to alter the implementation of dynamic linking in response to learners’ needs, fading this support as their knowledge and experience grows. Thus, when learners are new to the task, full linking could be provided between representations. As their experience grows, then full linking could be replaced by some signalling of the mapping between representations. Finally, if learners can make the representations reflect each other manually (acting as the dynamic linking did initially), then they should be able to work independently on either representation. None of the systems reported in this paper yet takes this approach to designing for deeper understanding, although many could be altered to adapt to this view.

5. Measuring learning with MERs

So far, it has been argued that MERs are used to support many different functions and that these functions can be distinguished by the role that translation plays in delivering these functions. This leaves software developers and teachers with a further important question — how can they tell when a multi-representational learning environment is successful.

As MERs are used for many different purposes, the learning objectives they are designed to support require different assessment. Again, given the varying roles of translation in the process, it is unsurprising that ways of assessing the successful learning differ in the need to identify whether learners can understand the relation between representations in addition to understanding each representation in isolation.

5.1. Complementary information and processes

When multi-representational learning environments are used for this purpose, it was argued that learners should not be required to understand the relation between the representations. Consequently, measures of performance taken to determine effectiveness of teaching do not require translation. For example, if each representation in the learning environment provides different information such that no redundancy of information exists between the representations in the system, then it is logically fairly simple to determine if each representation has been mastered. In this case, one would expect to see improvement in performance which required those dimensions of information that were presented in an understood representation and little or no improvement if a learner had not mastered the representation. For example, the representations used in CENTS can separately express the direction and magnitude of the accuracy of an estimate. If learners' performance only improved on the direction representation, then when learning outcomes were examined, one would expect to see little improvement at understanding magnitude but better understanding of direction.

When MERs are used to support different computational properties, then we need to assess competence on each representation in isolation. For example, we know that different strategies are associated with representations that have different computational properties (e.g. Tabachneck et al., 1994). Examining the way that learners understand the syntax and semantics of each representation and the strategies it supports will allow us insight into the effectiveness of the multi-representational learning environment.

5.2. Constraining interpretation

When MERs are used to constrain interpretation, a second representation is often designed to support interpretation of an unfamiliar representation and does not itself provide new information. Learners are not expected to construct an understanding of the set of relations between representations; rather, they are required to exploit these relations to understand a more complex representation. This principle informs the type of learning outcome that designers and teachers should be exploring. For example, the Skater World and DM³ microworlds described above present a simulation alongside representations such as a velocity–time graph. If learners can see the relation between the two, then we would expect that misreadings of the velocity–time graph (e.g. that a horizontal line on a velocity–time graph represents a stationary object and that negative gradient must entail negative direction) would disappear. Hence, assessment of the success of this use of MERs can again be measured by determining performance on representations in isolation. In this case, we need to identify whether their understanding of the constrained representation has improved as a result of using the learning environment.

5.3. Deeper understanding

If MERs are to be used to encourage deeper understanding by abstraction or extension, then learners must come to understand fully the relation between the representations. It is not

sufficient to measure performance on representations in isolation, in addition we need to understand whether learners can translate between them.

A number of techniques have been developed to explore whether learners can translate between representations. These include micro-genetic accounts (e.g. Schoenfeld et al., 1993) and computational modelling (Tabachneck-Schijf, Leonardo, & Simon, 1997). These methods while a useful tool for researcher are too time consuming to be applicable for formative evaluation or used by teachers in the classroom. One solution is to use learning environments that capture behavioural protocols. CENTS (see Fig. 1 described above) tracks how well learners understand the format and operators of each representation and domain knowledge. Furthermore, it also logs how well they understand the relation between representations. CENTS evaluates the similarity of learners' behaviour over the two representations. Performance on both representations is correlated to give a measure of how well learners understood the relation between representations (representational co-ordination). For example, the poorer performance of children in the mixed condition (Ainsworth et al., 1996) can be explained by examining this measure. Children in the mathematical and pictorial conditions showed a significant improvement in representational co-ordination over time, but there was no evidence that learners in the mixed condition understood the relation between the representations as they never converged their behaviour over the two representations. Furthermore, research with CENTS has shown that learning outcomes are not sufficient to determine whether translation has occurred. Ainsworth, Bibby and Wood (1997) showed that some learners could master aspects of estimation accuracy by focusing on a single representation rather than translating information between them. This knowledge is crucial if you wish to encourage abstraction or extension or have distributed information between representations; in each of these cases, any selective focusing on one representation will defeat the learning objective.

Measures such as representational co-ordination could be used by software developers to assess the design of learning environments during formative evaluation or by teachers (or intelligent tutoring systems) to monitor performance of individual children during learning sessions. They are appropriate when using MERs to encourage deeper understanding as they indicate the extent to which learners can see the relation between representations rather than focusing on how well they understand the domain. They can be used to predict learning outcomes and also could form the basis of dynamic models used to determine the degree of scaffolding learners need to understand the relation between representations.

6. Conclusions

This paper has presented a functional taxonomy of MERs, illustrated by reference to other systems described in this special issue. The variety of roles that MERs can play in supporting learning has been acknowledged. These roles are taken as the basis of effective design of multi-representational learning environments. They also determine the type of learning measures that designers, teachers and systems need to capture in order to determine whether users can learn successfully with the systems.

Differences in the design principles for each of class of MERs were defined in terms of

differences involved in the processes of translation between representations. If MERs are designed to support different information and processes then translation should be discouraged. If MERs are used to constrain interpretation, then translation should be automated. Finally, if MERs are used to develop deeper understanding, then translation should be scaffolded. These principles are speculative: as yet no research has examined the role of translation in learning environments in the light of these different claims. However, they serve as heuristics to guide further experimentation and serve to identify the kinds of learning measures that experimenters should collect. Such studies may then inform the design of the next generation of multi-representational learning environments.

Acknowledgements

This research was supported by the ESRC at the ESRC Centre for Research in Development Instruction and Training. I would like to thank Ben du Boulay, John Gardner, Rose Luckin and David Wood for commenting on an early version of this paper and Peter Cheng, Martin Oliver and Heather Wood for kindly providing me with pictures of their systems.

References

- Ainsworth, S. E., Wood, D. J., & Bibby, P. A. (1996). Co-ordinating multiple representations in computer based learning environments. In P. Brna, A. Paiva, & J. Self, *Proceedings of the European conference of artificial intelligence in education* (pp. 336–342). Lisbon: Edições Colibri.
- Ainsworth, S. E., Wood, D. J., & Bibby, P. A. (1997). Evaluating principles for multi-representational learning environments. In S. Vosniadou, K. Matsagouras, K. Maridaki-Kassotaki, & S. Kotsanis, *7th European conference for research on learning and instruction* (pp. 500–501). Athens: Gutenberg University Publications.
- Ainsworth, S. E., Wood, D. J., & Bibby, P. A. (1998a). Analysing the costs and benefits of multi-representational learning environments. In M. van Someren, H. P. A. Boshuizen, T. de Jong, & P. Reimann, *Learning with multiple representations* (pp. 120–134). Oxford: Elsevier Science.
- Ainsworth, S. E., Wood, D. J., & O'Malley, C. (1998b). There's more than one way to solve a problem: Evaluating a learning environment to support the development of children's multiplication skills. *Learning and Instruction*, 8(2), 141–157.
- Anderson, J. R. (1983). *The architecture of cognition*. Cambridge, MA: Harvard University Press.
- Bibby, P. A., & Payne, S. J. (1993). Internalization and the use specificity of device knowledge. *Human-Computer Interaction*, 8(1), 25–56.
- Cheng, P. C.-H. (1996). Law encoding diagrams for instructional systems. *Journal of Artificial Intelligence in Education*, 7(1), 33–74.
- Confrey J. 1992. Function Probe©. Santa Barbara, CA: Intellimation library for the Macintosh.
- Cox, R., & Brna, P. (1995). Supporting the use of external representations in problem solving: The need for flexible learning environments. *Journal of Artificial Intelligence in Education*, 6(2/3), 239–302.
- Cronbach, L. J., & Snow, R. E. (1977). *Aptitudes and instructional methods*. New York: Irvington.
- Dienes, Z. (1973). *The six stages in the process of learning mathematics*. Slough: NFER-Nelson.
- Dugdale, S. (1982). Green globs: A micro-computer application for graphing of equations. *Mathematics Teacher*, 75, 208–214.
- Ehrlich, K., & Johnson-Laird, P. N. (1982). Spatial descriptions and referential continuity. *Journal of Verbal Learning and Verbal Behaviour*, 21, 296–306.

- Gilmore, D. J., & Green, T. R. G. (1984). Comprehension and recall of miniature programs. *International Journal of Man-Machine Studies*, 21, 31–48.
- Giunchiglia, F., & Walsh, T. (1992). A theory of abstraction. *Artificial Intelligence*, 57, 323–389.
- Hennessy, S., Twigger, D., Driver, R., O'Shea, T., O'Malley, C., Byard, M., Draper, S., Hartley, R., Mohamed, R., & Scanlon, E. (1995). Design of a computer-augmented curriculum for mechanics. *International Journal of Science Education*, 17(1), 75–92.
- Kaput, J. J. (1989). Linking representations in the symbol systems of algebra. In S. Wagner, & C. Kieran, *Research issues in the learning and teaching of algebra* (pp. 167–194). Hillsdale, NJ: Erlbaum.
- Kaput, J. J. (1992). Technology and mathematics education. In D. A. Grouws, *Handbook of teaching and learning mathematics* (pp. 515–556). New York: Macmillan.
- Larkin, J. H., & Simon, H. A. (1987). Why a diagram is (sometimes) worth ten thousand words. *Cognitive Science*, 11, 65–99.
- Meyer, J., Shinar, D., & Leiser, D. (1997). Multiple factors that determine performance with tables and graphs. *Human Factors*, 39(2), 268–286.
- Oliver M. 1998. Visualisation and manipulation tools for modal logic. Unpublished Ph.D. thesis, Open University.
- Oliver, M., & O'Shea, T. (1996). Using the model-view-controller mechanism to combine representations of possible worlds for learners of modal logic. In P. Brna, A. Paiva, & J. Self, *Proceedings of the European conference of artificial intelligence in education* (pp. 357–363). Lisbon: Edições Colibri.
- Pheasey, K., O'Malley, C., & Ding, S. (1997). A distributed collaborative micro-world to teach Newtonian mechanics. In S. Vosniadou, K. Matsagouras, K. Maridaki-Kassotaki, & S. Kotsanis, *7th European conference for research on learning and instruction*. Athens: Gutenberg University Publications.
- Resnick, L. (1992). From protoquantities to operators: Building mathematical competence on a foundation of everyday knowledge. In G. Leinhardt, R. Putnam, & R. A. Hatrup, *Analysis of arithmetic for mathematics education* (pp. 375–429). Hillsdale, NJ: Erlbaum.
- Resnick, L., & Omanson, S. (1987). Learning to understand arithmetic. In R. Glaser, *Advances in instructional psychology volume 3* (pp. 41–95). Hillsdale, NJ: LEA.
- Roberts, M. J., Wood, D. J., & Gilmore, D. J. (1994). The sentence-picture verification task: Methodological and theoretical difficulties. *British Journal of Psychology*, 85, 413–432.
- Scaife, M., & Rogers, Y. (1996). External cognition: how do graphical representations work? *International Journal of Human-Computer Studies*, 45, 185–213.
- Schoenfeld, A. (1986). On having and using geometric knowledge. In J. Hiebert, *Conceptual and procedural knowledge: The case of mathematics* (pp. 225–265). Hillsdale, NJ: Erlbaum.
- Schoenfeld, A. H., Smith, J. P., & Arcavi, A. (1993). Learning: The microgenetic analysis of one student's evolving understanding of a complex subject matter domain. In R. Glaser, *Advances in instructional psychology volume 4* (pp. 55–175). Hillside, NJ: Erlbaum.
- Schwartz, D. L. (1995). The emergence of abstract representations in dyad problem solving. *The Journal of the Learning Sciences*, 4(3), 321–354.
- Sfard, A. (1991). On the dual nature of mathematical conceptions: reflections on processes and objects as different sides of the same coin. *Educational Studies in Mathematics*, 22, 1–36.
- Snow, R. E., & Yalow, E. (1982). Education and intelligence. In R. J. Sternberg, *A handbook of human intelligence* (pp. 493–586). Cambridge: Cambridge University Press.
- 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., Koedinger, K. R., & Nathan, M. J. (1994). Towards a theoretical account of strategy use and sense making in mathematical problem solving. In A. Ram, & K. Eiselt, *Proceedings of the 16th annual conference of the cognitive science society* (pp. 836–841). Hillsdale, NJ: Erlbaum.
- Tabachneck, H. J. M., Leonardo, A. M., & Simon, H. A. (1994). How does an expert use a graph? A model of visual & verbal inferencing in economics. In A. Ram, & K. Eiselt, *Proceedings of the 16th annual conference of the cognitive science society* (pp. 842–847). Hillsdale, NJ: Erlbaum.
- Tabachneck-Schijf, H. J. M., Leonardo, A. M., & Simon, H. A. (1997). CaMeRa: A computational model of multiple representations. *Cognitive Science*, 21(3), 305–350.

- Thompson, P. W. (1992). Notations, conventions and constraints: Contributions to effective uses of concrete materials in elementary mathematics. *Journal for Research in Mathematics Education*, 23(2), 123–147.
- Underwood, J., & Underwood, G. (1987). Data organisation and retrieval by children. *British Journal of Educational Psychology*, 57, 313–329.
- Winn, B. (1987). Charts, graphs and diagrams in educational materials. In D. M. Willows, & H. A. Houghton, *The psychology of illustration* (pp. 152–198). New York: Springer.
- Wood, D. J., Bruner, J. S., & Ross, G. (1976). The role of tutoring in problem solving. *Journal of Child Psychology and Psychiatry*, 17, 89–100.
- Yerushalmy, M. (1991). Student perceptions of aspects of algebraic function using multiple representation software. *Journal of Computer Assisted Learning*, 7, 42–57.