

DESIGNING EFFECTIVE MULTI- REPRESENTATIONAL LEARNING ENVIRONMENTS

Shaaron Ainsworth

*ESRC Centre for Research in Development, Instruction &
Training Department of Psychology
University of Nottingham
University Park
Nottingham NG7 2RD
United Kingdom*

email: Shaaron.Ainsworth @nottingham.ac.uk

March 1999

For learning with multiple external representations (MERS) to be successful, the design of a learning environment must take advantage of the properties of different representations without overwhelming a learner with their associated costs. This paper presents an analytic framework that consists of a description of the functions of MERS, an analysis of the learning demands of using MERS and consideration of the design decisions that uniquely apply to multi-representational learning environments. These are integrated to propose a set of idealised designs for each function of MERS. This framework was constructed for two purposes. Firstly, it can be used to compare existing learning environments and so allow more accurate generalisations from previous empirical work. Secondly, it is intended to provide the basis for further experimentation in order to develop effective design principles for multi-representational learning.

TABLE OF CONTENTS

1.0 Introduction	1
2.0 Functions Of Multiple Representations	2
2.1 Different Information And Different Processes	2
2.2 Constraints On Interpretation	7
2.3 Deeper Understanding	9
3.0 Learning Demands Of Multiple Representations	15
3.1 Learning The Format And Operators Of A Representation	15
3.2 Learning The Relation Between The Domain And The Representation	16
3.3 Learning The Relation Between The Representations	17
3.4 Translating Between Multiple Representations	18
4.0 Design Decisions	25
4.1 Distributing Information Between Representations	27
4.2 Similarity Between Representations	32
4.3 Automatic Translation	40
4.4 Number Of Representations	43
4.5 Ordering And Sequencing Representations	44
5.0 Designing For Different Functions Of Multiple Representations	47
5. 1 Designing For Different Information And Processes	47
5.2 Designing For Constraining Interpretation	52
5.3 Designing For Deeper Understanding	56
5.4 Summary	60
6.0 Conclusion	60
References	61
Appendix One	70
Appendix Two	74

DESIGNING EFFECTIVE MULTI-REPRESENTATIONAL LEARNING ENVIRONMENTS

1.0 INTRODUCTION

Multi-representational learning environments are used by a wide range of learners in a number of domains and many advantages are claimed for their use. By using multiple external representations (MERs), it is hoped that learners can benefit from the properties of each of the representations and that ultimately this will lead to a deeper understanding of the subject being taught. However, research that has evaluated how effectively multi-representational environments support learning has produced mixed results. A number of studies have shown that learners find working with MERs to be very difficult (*e.g.*, Tabachneck, Leonardo & Simon, 1994; Yerushalmy, 1991). Consequently, designers of multi-representational learning environments are faced with the question of how to develop a system where the learners can benefit from the advantages of MERs without succumbing to their disadvantages.

There is abundant evidence of the important roles that external representations play in supporting learning (*e.g.* White, 1993; Zhang & Norman, 1994). Research on multiple external representations investigates the effects of combining different external representations upon learning. Consequently, one important criteria for clear research in this area is to define precisely different representations. Throughout this paper, descriptions of representations will be based upon Palmer's analysis (Palmer, 1978). He proposes that any particular representation should be described in terms of (1) the represented world, (2) the representing world, (3) what aspects of the represented world are being represented, (4) what aspects of the representing world are doing the modelling and (5) the correspondence between the two worlds. Using this definition of a representation, multi-representational systems will be considered in terms of the represented world and the representing world.

This paper presents a framework for addressing how to design effective multi-representational learning environments. It is composed of four parts. The first three sections set out the elements of the framework: a description of the various functions that MERs can play in learning environments, analysis of the learning demands of MERs and identification of the design decisions that are unique to multi-representational learning environments. In the final section, these three elements are combined to propose design guidelines for multi-representational software. These consist of a series of idealised designs for each function of MERs aimed at minimising the learning demands faced by the user. This should help ensure that users of multi-representational learning environments can benefit from the many advantages MERs bring to learning.

2.0 FUNCTIONS OF MULTIPLE REPRESENTATIONS

In order to design an effective multi-representational learning environment, the first question that should be considered is why use more than one representation. Generalised principles for effective learning with MERs will not occur until the variety of functions that MERs serve is recognised. In this section, a functional taxonomy of MERs will be proposed and illustrated.

Analysis of existing multi-representational environments suggests that there are three main functions that MERs serve in learning situations. The first function is to use representations that contain different information or provide different computational properties. In the second case, MERs are used to constrain possible (mis)interpretations of a representation or domain. Finally, MERs can be used to encourage deeper understanding of a situation. Each of these uses of MERs have several subclasses. Often a single multi-representational environment will be required to serve several of these functions, but to begin with, each will be considered separately.

2.1 Different Information and Different Processes

The first use of MERs is to combine representations that differ either in the information each expresses (the represented world) or in the processes each supports (the representing world). By combining representations that differ in these ways, it is hoped that learners will benefit from the advantages of each of the individual representations in the learning environment.

2.1.1 Using MERs to convey different information

One reason to use MERs is to vary the information that is expressed by each representations so that each representation denotes different aspects of the represented world. Multiple representations tend to be used for this purpose when a single representation would be insufficient to carry all the information about the domain or would be too complicated for learners to interpret if it did so.

‘MoLE’, Oliver & 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 Figure 1). The second is a concrete representation of a grid of polygons that illustrates the content of each world (RHS in Figure 1). In this case, there is no redundancy between the two representations as each expresses different information. A similar use of MERs is used within the Internet Software Visualisation

^{*} Appendix Two contains a short description of the most frequently discussed systems in this paper. It should be noted that summaries of the systems’ goals and representations reflect my own views of the systems and may not completely accord with the original author’s views. The desire to minimise such discord encouraged the use of systems to which I had access, rather than focusing on the better known systems.

Laboratory (Mulholland & Domaigne, 1997). The first representation shows the search path that a PROLOG interpreter takes when satisfying a subgoal and a second textual representation describes the detail of each predicate in turn.

It is possible that in both of these systems, one representation could have been created that carried all the information. For example, in MoLe the modal relation representation could have included the content grid within each node. Yet, had it done so, this representation would have quickly become cluttered and would have been difficult to interpret when more than a few worlds were displayed. Using MERs allowed the designers to create representations that are more readable. An analysis of the work scratchings suggests that dividing the information in this way might also have allowed learners to concentrate on different aspects of the task and made the learning goals more manageable (Oliver, 1998). This issue is considered in section 4.1 where the issue of designing for information redundancy in multi-representational learning environments is discussed.




Figure 1. The relation and world descriptions representations in MoLe

2.1.2 Using MERs to support new inferences by providing partially redundant representations

One specific case of distributing information over representations is where each of the representations describe different aspects of the represented world but maintain certain elements in common. This partial redundancy of information makes possible new interpretations about a domain. For example, one picture may provide the information that John is taller than Jill and the second that Jill is taller than Jack. By reasoning about the conjunction of these representations, a further inference can be drawn - that John is taller than Jack. In this example, the representations were of the same format. Yet, this is not a crucial factor as sometimes the representations may be of different formats. It is the partial

redundancy of information between the representations that is the defining characteristic of this function of multiple representations.

Another classic illustration of this situation is the problem of finding the quickest route between two London Underground stations. The London Underground map designed by Harry Beck in the 1930s does not preserve geographical and topological information. Its purpose is to represent connections between stations. By looking at this Underground map, it is a simple task to determine the shortest train journey between two stations. However, this does not guarantee that you will find the shortest and quickest route. A train journey between two stations that requires a number of changes could be reached on foot in a matter of minutes (*e.g.* Bank to Mansion House - a total of seven stations and two different lines or 200 metres by foot). An accurate solution to this problem could be found by integrating the information provided by a street map which preserves 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. However, one representation would often be too complex to interpret if it did so. By distributing information over partially redundant representations, multi-representational learning environments can use less complicated representations. However, learners are then faced with the task of integrating the information from these representations - a problem that is considered in detail in section 3.

2.1.3 Using MERs with different processes

A further use of MERs is to exploit the varying computational processes supported by different representations. For example, Larkin & 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 information explicit, emphasise empty cells (thus directing attention to unexplored alternatives), allow quicker and more accurate readoff and highlight patterns and regularities. 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. Graphs show trends and interaction more successfully than alphanumeric representations. Given these conclusions, it is not surprising that one of the most common reason to use MERs in learning environments is to obtain the different computational advantages of each of the individual representations.




Figure 2. ReMIS-CL

One system that uses a number of representations to support different processes is ReMIS-CL (Cheng, 1996a). Seven different representations are available to help students understand the nature of elastic collisions. Many of these representations are Law Encoding Diagrams (LEDs). An LED is a representation that correctly encodes the underlying relations of a law(s) by means of geometric, topological or spatial constraints such that each instantiation of a LED represents one instance of the phenomenon or one case of the laws. The representations include numerical/equations, one dimensional property diagrams, mass velocity diagrams, velocity-velocity graphs. A learner/instructor can choose up to four of these representations to view simultaneously. In addition, an animation of the collision is always available. The one dimensional property diagram (LHS of Figure 2) makes explicit the structure of the situation. The lines that give the initial and final velocities can effectively be overlaid on the simulation. In contrast, the velocity-velocity graph (RHS of Figure 2) emphasises that velocities are given by the simultaneous satisfaction of two separate relations. The two intersections between the diagonal and the ellipse indicate that two pairs of values can be found. Empirical support for the value of these representations compared to the traditional equations can be found in Cheng (1996b).

There is a large body of literature concerning how these computational properties of different representations influence learning and problem solving. These effects can be shown to be advantageous at task, strategy and learner levels.

To achieve a particular objective, a learner is normally required to perform a number of different tasks. Yet, there is rarely a single representation that is absolutely good, rather particular representations facilitate performance on certain tasks. This point was made by

Gilmore & Green (1984) who proposed the match-mismatch conjecture - that performance would be facilitated when the form of information required by the problem matches the form provided by the notation. This analysis has subsequently been applied to a number of domains (*e.g.* comparing visual and textual programming languages; Green, Bellamy & Petre, 1991). Empirical support for this conjecture is provided by Bibby and Payne (1993) who gave subjects instructions on how to operate a simple control panel device using either (informationally equivalent) tables, procedures, diagrams. To learn to operate the device fully, a number of different tasks needed to be performed. These include detecting faulty components and altering switch positions. No single representation was better overall, but there were significant interactions between task and representation. Subjects given tables and diagrams identified faulty components faster, but those given procedures were faster at deciding which switches were mispositioned. Providing learners with MERs may in the short term decrease performance as they have more representations to understand, but, in the longer term may facilitate understanding as learners will have the opportunity to apply the most appropriate representation to solve different aspects of the task.

Different representations have been shown to promote different strategies. For example, Tabachneck, Koedinger & Nathan (1994) examined learners solving algebra word problems. They identified six external representations (*e.g.* verbal arithmetic, diagrams and written algebra) which were associated with four strategies (algebra, guess-and-test, verbal-math and diagram). No single strategy was more effective than any other, but the use of multiple strategies was about twice as effective as any strategy used alone. As each strategy had inherent weaknesses, switching between strategies made problem solving more successful by compensating for this. Cox (1996) observed a similar effect when students solved analytical reasoning problems. He found students used a variety of representations (*e.g.* logic, set diagrams, tables, and natural language). In 17% of cases, subjects used more than one representation and this tended to be associated with good performance. Both Cox and Tabachneck found that it was at impasses that subjects tended to switch between representations. Consequently, where learners are given the opportunity to use MERs, they may be able to compensate for weaknesses associated with one particular strategy and representation by switching to another.

A third explanation often provided for this use of MERs is that there are individual differences in representational and strategic preference. If alternative representations are provided, users can act upon the representation of their choice. Research examining the impact of various personality or cognitive factors in relation to learning with external representations has proposed differential effects of, *inter alia*, IQ, spatial reasoning, locus of control, field dependence, verbal ability, vocabulary, gender and age (see Winn, 1987). A common (although by no means consistent finding) is that lower ability learners benefit from graphical representations of the task (see Cronbach & Snow, 1977; Snow & Yalow, 1982).

Cognitive style is a somewhat contentious issue with noted intra-individual differences as well as inter-individual differences and there is not necessarily a simple relation between preferred style and task performance (*e.g.* Roberts, Wood, & Gilmore, 1994). However, an account based on the premise that learners will often have varying experience and expertise with different representations would also suggest that it can be beneficial to provide learners with alternative representations.

It can be seen that there may be considerable advantages for learning with MERs that provide different information or support different inferences. By combining representations, learners are no longer limited by the strengths and weaknesses of one particular representation.

2.2 Constraints on Interpretation

A second use of MERs is to help learners develop a better understanding of a domain by constraining their interpretation of the representations and tasks. This can be achieved in two ways: (a) by employing a familiar or concrete representation to support the interpretation of a second abstract or unfamiliar representation and; (b) by exploiting inherent properties of a representation to constrain interpretation of a second representation.

2.2.1 Using MERs so that a familiar or concrete representation constrains interpretation of a second unfamiliar or abstract representation

In this case, an additional representation may be employed to support the interpretation of a more complicated, abstract or unfamiliar representation. Therefore, the second representation is used to provide support for a learner's missing or erroneous knowledge. Commonly simulation environments exploit MERs to this end. For example, microworlds such as DM³ (Hennessy *et al.*, 1995) or SkaterWorld (Pheasey, Ding & O'Malley, 1997; see Figure 5) provide a simulation of a skater alongside a velocity-time graph (amongst other representations). Two of the misconceptions common to children learning Newtonian mechanics are that a horizontal line on a velocity-time graph must represent a stationary object and that negative gradient must entail negative direction. When the simulation shows the skater still moving forward, students are given the opportunity to debug their misinterpretations of the line-graph. ReMIS-CL (Cheng, 1996a) teaches the physics of elastic collisions. A novel class of representation, Law Encoding Diagrams (LEDs), are presented for learners to reason with and act upon. User's reasoning about information presented in the LEDs (*e.g.* initial and final velocities) can be debugged with comparison to an animated simulation of the collision (Figure 2).

Multimedia systems often exploit this aspect of MERs (*e.g.* Millwood, 1996), for example, by providing written and spoken text simultaneously. If children are developing reading skills and find the written text difficult, or if the spoken text is hard to understand

(*e.g.* Shakespearean language or speech with a broad regional accent), then presence of the constraining representation may help support understanding of the first representation.

The primary purpose of the constraining representation in all of these examples is **not** to provide new information but to support a learner's reasoning about a second representation. 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 constrains interpretation of a second representation

In contrast, sometimes the more abstract or unfamiliar representation is used to constrain interpretation of the second representation. In this case, it does so by exploiting an inherent property of the representation. For example, it is argued that graphical representations are less expressive than many propositional representations (*e.g.* Stenning & Oberlander, 1995). This can be seen in the ambiguity permitted in the propositional representation 'the knife is beside the fork' which is completely permissible. However, an equivalent image would have to picture the fork as either to the left or to the right of the knife (*e.g.* Erhlich & Johnson-Laird, 1982). Thus, when these two representations are presented together, interpretation of the first representation may be constrained when the representational system is considered as a whole.

This function of MERs can be seen in the design of multi-representational learning environments. For example, COPPERS (Ainsworth, Wood & O'Malley, 1998) teaches children about multiple solutions to coin problems. Two representations are used to describe each of the children's solution in detail (see Figure 3). The first one is a familiar place value representation (RHS Figure 3). The user is reminded of how many of each type of coin they used in such a way as to make explicit the arithmetic operations. In the second, a more unfamiliar tabular representation expresses the same information (per single row), but the operations are implicit in the values in the cells and column headings (LHS, Figure 3). The main role of the place value representation is to constrain the possible misinterpretations of the unfamiliar table representation by indicating the appropriate format and operators for the table representation (the first type of constraint). However, the tabular representation can in turn constrain interpretation of the place value representation (the second type of constraint). Answers to coin problems such as '5p, 10p, 5p, 10p' and '5p, 5p 10p, 10p' may appear very different to young children if they do not understand commutativity. The tabular representation of coin values used in COPPERS does not express ordering information. Therefore, if children translate between the representations, the equivalence of the two different orderings in the place value representation is more likely to be recognised.

PREVIOUS								ANSWERS	
1p	2p	5p	10p	20p	50p	£1	TOTAL		
			3	2	1		70p	Good	
		4	3	2			70p	one of the answers to this	
		14	3	1			70p	problem is	
5	5	3	2	1			70p	2 x 5 pence =	10p +
2	4	2	1	2			70p	1 x 10 pence =	10p +
							70p	2 x 20 pence =	40p +
							70p	2 x 1 pence =	2p +
								4 x 2 pence =	8p
									<hr/>
									70p
									<hr/>



 NEXT

Figure 3. COPPERS - place value feedback and the summary table

Therefore, a further function for MERs is to constrain interpretation either by supporting missing knowledge or through providing representations that encourage different interpretations of the situations.

2.3 Deeper Understanding

It has also been claimed that multiple representations can lead 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 section, ‘deeper understanding’ will be considered in terms of using MERs to promote abstraction, to encourage generalisation and to teach the relation between representations. The differences between these functions of MERs are quite subtle and all may be present at some stage in the life cycle of encouraging deeper understanding with a multi-representational environment.

2.3.1 Using MERs to support abstraction

Abstraction has been defined in a number of different ways. One common sense of the term is as ‘subtraction’ where the emphasis is on extracting only a portion of original representation. For example, Giunchiglia & Walsh (1992) in defining abstraction refer to ‘throwing away details’. More completely, they (informally) define abstraction as:

1. The process of mapping... the ground representation onto a new representations called the abstract representation which:
2. helps deal with the problem in the original search space by preserving certain desirable properties and

3. is simpler to handle as it is constructed from the ground representation by throwing away details

It should be noted that this definition derives from an artificial intelligence perspective and no psychological claim should necessarily be implied about whether humans find abstract representations simpler to handle.

An alternative conceptualisation of abstraction emphasises re-ontologisation. For example, 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 resulting 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 reading emphasising either an operational or a structural view. Such an analysis is not only appropriate for school taught subjects – a child may learn to count **on** their fingers and then begin to count other objects **with** their fingers. Whether abstraction occurs by subtraction, reification or by re-ontologisation, there is a common sense that the abstracted understanding that results is somehow ‘higher’ than the original representations.

So how might multiple representations encourage abstraction? It is hoped that if you provide learners with a rich source of representations of a domain, then they will build references across these representations. Such knowledge can then be used to expose the underlying structure of the domain represented. For example, Dienes (*e.g.* Dienes 1973) argues that perceptual variability (the same concepts represented in varying ways) provides learners with the opportunities for building such abstractions.

Resnick & Omanson (1987) and Schoenfeld (1986) describe the process of learning to add and subtract with Dienes blocks and written numerals. During a substantial intervention program, children were given mapping instructions about the correspondence of these two representational systems. If children understood the operations on the Dienes blocks and the mapping between these concrete manipulatives and the symbolic procedure, then they should master the symbolic procedures. In the sense of abstraction defined above this is not an abstracted understanding as no higher-level knowledge results. However, an alternative conceptualisation is that mapping instruction could have taught children to identify parallel structures in two symbolic domains (trading with Dienes blocks, carrying, and borrowing with base 10 algorithms). Having done this, children could then see that the structure of base 10 arithmetic in an abstracted sense by recognising that these are both actions on quantities. Although this particular intervention did not lead to many of the children requiring an abstract understanding of subtraction, base 10 and number representations, it does suggest a route by which multiple representations can serve such a goal

Schwartz (1995) provides interesting converging evidence that multiple representations can generate more abstract understanding. In this case, the multiple representations are

provided by different members of a collaborating pair. With a number of 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.

Although there is some evidence that multiple representations can lead children to a more abstract representation, little is known about how to design for abstraction or the conditions under which abstraction might be beneficial. This issue is considered further in section 5.3.

2.3.2 Using MERs to support extension

Extension or generalisation can be considered as a way of extending knowledge that a learner 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. For example, an extension of the concept of triangle from red objects with three sides whose internal angles add up 180 degrees to blue objects with three sides whose internal angles add up 180 degrees is generalisation. If a child says any coloured object with three sides whose internal angles add up 180 degrees is a triangle then this new definition is still extension of the old one. In the definition of abstraction given in this paper, a child's concept of triangle would not be considered as abstracted until he or she realises that they need not refer to colour at all and so change their rule to objects with three sides whose internal angles add up 180 degrees are triangles. In cognitive models such as ACT*, generalisation often occurs through variablisation (*e.g.* Anderson, 1983).

When considering representations, extension can refer to two different aspects of a learning situation - extending the domains where a representation is used or extending the way that domain knowledge is embodied to include another representation.




Figure 4. Extending knowledge to (a) new domains and (b) new representations*

The first case of extension can be seen whenever a representation, taught for one purpose or domain, is used to serve another (LHS of Figure 4). 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.* The problem in this case is encouraging learners to apply representations outside the initial context given the context sensitive nature of learning. However, for the purposes of this paper where the focus is on multiple representations, this type of extension although common in learning situations, is not strictly relevant as it describes a single representation in multiple domains.

The second type of extension is extending domain knowledge through a variety of representations (RHS of Figure 4). 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 that knowledge to see acceleration in such representations as tables, acceleration-time graphs, tickertape *etc.* This process can be considered extension if a learner proceeds from understanding how one representation expresses the concept to understanding how a second representation can embody the same knowledge. Using MERs for this purpose is quite close to that of constraining interpretation. However, it differs in the emphasis placed on understanding the relation between two representations. When supporting extension with MERs, the emphasis is placed on teaching children how their existing knowledge can be extended to new representations. In contrast, when constraining interpretation between representations the intention is to exploit knowledge of the relation between two representations to some further end.

2.3.3 Using MERs to teach relation among representations.

* Dark lines are used to refer to translation processes and dotted lines to the relationship between each representation and the domain. The length of the lines intended to indicate the amount of work required to map between representations or between a representation and a domain.

This function of MERs is only subtly different from the cases we have already considered. Similarly to extension, the pedagogical goal is to teach learners 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 (see section 2.2.1). Other representations include tickertape, force arrows, net force indicator, tables of velocity, distance travelled and time elapsed (Figure 5). In addition, learners can choose one from velocity-time, distance-time or acceleration-time graphs. In experiments with the system, it could be seen that learners in this domain rarely have a full understanding of any single representation. Much of the learning that takes place with this system is directed at relating these different representations.

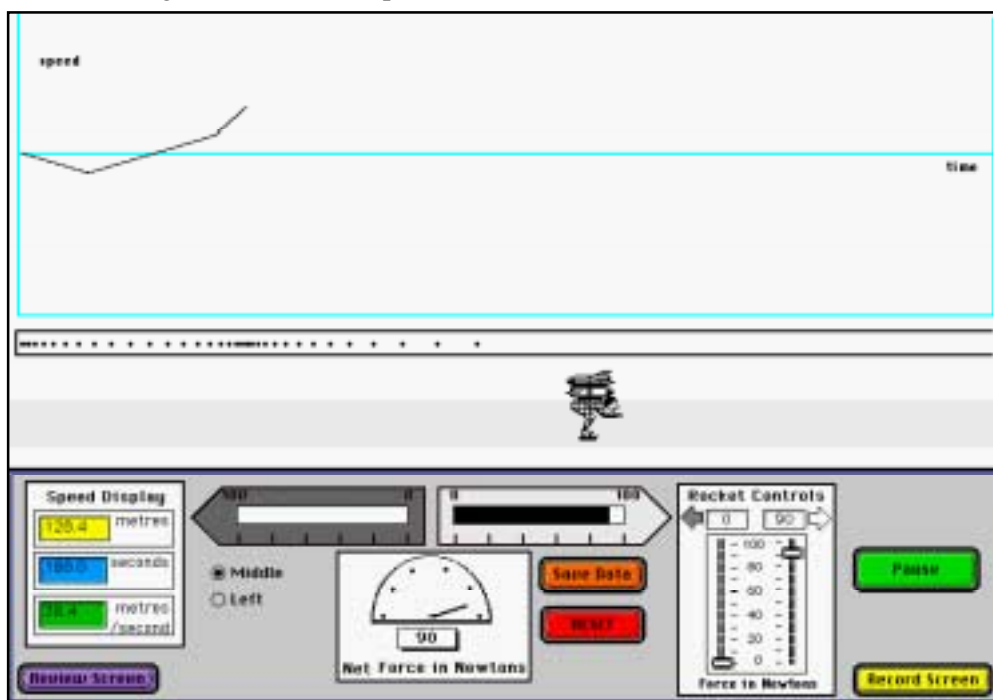


Figure 5. SkaterWorld showing the simulation screen

The QUADRATIC Tutor (Wood & Wood, in press) teaches pupils with only limited experience of algebra to develop an understanding of the quadratic function. In particular, it uses the area of squares to make salient the properties of algebraic expressions. QUADRATIC is designed to teach children about the equivalences of the geometric and algebraic representations. Learners can come to understand that $x^2 + 2x + 1 = (x+1)^2$ by referring to a graphic representation of the $x+1$ square, (see Figure 6).

Teaching progresses by allowing children to construct squares of different sizes, then to relate the algebraic expressions to a diagram and to expand the general case. This is then repeated for the $(x+n)^3$ and the $(x-n)^2$ cases. The fairly 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 unfamiliar representations. However, if learners already have substantial knowledge of algebra, then by using QUADRATIC they could extend this knowledge to the novel situation of explaining the properties of square




Figure 6. The Quadratic Tutor

The goal of teaching relation between representations can sometimes be an end in itself. For example, much emphasis is placed on learning how to construct a graph given an equation (*e.g.* Dugdale, 1982). However, 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. For example, Schoenfeld (1986) suggests that an initial characterisation of using Dienes blocks to support the understanding of the symbolic procedures by teaching the mapping between the concrete manipulative and the symbolic procedures is better understood as requiring abstraction (see section 2.3.1). 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 learners master these processes that their knowledge of the representations can serve other ends.

2.4 Summary

There are many different reasons why MERs can be beneficial for learning. Research was reviewed and it was suggested that MERs are commonly used for one of three main purposes (*i.e.* that MERs support different ideas and processes, can constrain interpretations and promote 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. For example, representations used because of their different computational properties may also encourage abstraction if learners can map over them. However, for these objectives to be met, learners must meet a number of significant learning demands. These are discussed in the next section.

3.0 LEARNING DEMANDS OF MULTIPLE REPRESENTATIONS

Learners are faced with complex learning tasks when they are **presented** with a novel multi-representational system. They must learn the format and operators of each representation, understand the relation between each representation and the domain it represents and learn how the representations relate to each other. A fourth source of learning demand is only present when learners must construct or select their own representations. This is discussed in section 4.2. The following section will give examples of each of these learning demands and the problems associated with them. As the task of translating between representations is unique to multi-representational system, this learning demand will be discussed in more detail. These cognitive tasks are presented in sequence but it should not be inferred from this that learners will approach the task of understanding a multi-representational system in this same order.

3.1 Learning the Format and Operators of a Representation

The first learning task facing any user of a representation is to understand each representation. They must know how a representation encodes and presents information (the 'format'). In the case of a graph, the format would be attributes such as lines, labels, and axes. They must also learn what the 'operators' are for a given representation. For a graph, operators to be learnt include how to find the gradients of lines, maxima and minima, intercepts, *etc.* At least initially, such learning demands will be great, and will obviously increase with the number of representations employed.

A number of studies have shown the difficulties that learners face in understanding this aspect of representations. Preece (1983) reports that 14-15 year old children experienced difficulty in applying and understanding the format and operators of graphs. For example, some pupils have trouble with reading and plotting points, they interpreted intervals as points, confused gradients with maxima and minima, *etc.* Petre (1993) describes some similar effects when adults are learning to understand a visual interface (countering the familiar claim that graphical representations are inherently better than textual ones as they require no learning in order to use them). In observing differences between novices and experts, she

showed that novices lack proficiency in secondary notation (*i.e.* perceptual cues that are not described by the formal semantics of a representation). Novices may find navigation of graphical representations difficult as they don't have the required reading and search strategies and in contrast to expert performance, they tend not to match strategies to the available representations. Additionally, the operators of one representation are often used inappropriately to interpret a different representation. A representation of graph may be interpreted using the operators for pictures. This behaviour is seen when learners are given a velocity-time graph of a cyclist travelling over a hill. They should select a U shaped graph, yet many show a preference for graphs with a hill shaped curve (*e.g.* Kaput, 1989).

3.2 Learning the Relation between the Domain and the Representation

Learners must also come to understand the relation between the representation and the domain it is representing. This task will be particularly difficult for learning with MERs as opposed to problem solving or reasoning, as learners will also have incomplete domain knowledge. Learners must know which operators to apply to the representation to retrieve the relevant domain information. To return to the graph example, children must learn when it is appropriate to examine the slope of a line, the height of a line, or the area under a line. For example, when attempting to read the velocity of an object from a distance-time graph, children often examine the height of line, rather than the gradient.

Brna (1996) provides details from a number of domains about the difficulties learners face when attempting to relate a representation to a domain. For example, even fairly competent programmers who had received information about the elements of a new (visual programming) representation failed to clearly map the format of the new representation onto their existing domain knowledge. Laborde (1996) discusses the difficulties that students had in connecting geometrical properties to spatial properties when learning with Cabri-géomètre. Encouragingly, though, she believes that the computer environment acted to help children learn these relations by enlarging the range of visual phenomena possible (for example by dragging circles, tangents, *etc.*) whilst at the same time constructing these visualisations in a theoretically meaningful way. These problems do not only arise with abstract representation such as graphs, visual programming languages or geometric objects. Boulton-Lewis & Halford (1990) point out that even concrete representation such as Dienes blocks and fingers still need to be mapped to domain knowledge. Processing loads may still be too high for children to obtain the anticipated benefits of such apparently simple representations.

3.3 Learning the Relation between the Representations

When MERs are presented together, learners must come to understand how representations relate to each other. This task is unique to multi-representational situations. A number of researchers have noted the problems that novices have in learning the relation between representations. Tabachneck, Leonardo & Simon (1994) report that novices learning with

MERs in economics did not attempt to integrate information between line graphs and written information. Students' performance on quantitative problems where answers could be read off from graphs was good, but it was poor on problems requiring explanation and justification. A similar pattern of results was found for graph generation as well as interpretation. This contrasted with expert performance where graphical and verbal explanations were tied closely together. Similarly, Yerushamly (1991) examined 35 fourteen-year-olds understanding of functions after an intensive three-month course with multi-representational software. He found that only 12% of students gave answers that involved both visual and numerical considerations. Lesh, Post & Behr (1987) provide off-line examples of the difficulties that children have in translating between representations. In an apparently simple problem of choosing which of three pictures showed 1/3rd shaded, grade school pupils' and even college students' performance was surprisingly poor. For example, only 25% of 12 to 13-year-old children could select the right answer.

Competent performance with MERs requires learners to notice both regularities and discrepancies between representations (Borba, 1994; Confrey, 1994). Yet, Yerushamly (1991) found that students seemed unaware of contradictions between answers in the different representations. DuFour-Janvier, Bednarz & Belanger (1987) report a similar phenomenon. When children were asked to subtract using both an abacus and conventional written symbols, they commonly did not recognise the correspondence between the two representations and were unconcerned if they obtained different answers from each representation.

Three different learning demands of presented MERs have been described. It is obvious from this discussion that learners will not be able to benefit fully from the proposed advantages of MERs if they cannot meet these demands. Each time a new representation is introduced to a multi-representational system, these demands increase. In all cases, the format and operators of a representation must be understood (at least to some degree) as must the relationship between the representation and the domain. In addition, as translation between the different representations is required for many of the uses of MERs, increases in learning demands will not be simply additive. Studies in many domains have shown how difficult translating between representations can be for learners. In the next section, the variety of ways that translation between representations have been studied are presented.

3.4 Translating Between Multiple Representations

There is considerable evidence that learners find translating between representations difficult. They frequently do not use more than one representation, even after extensive training with multi-representational software. Even when they are required to do so, they seem to treat each representation in isolation, not noticing the regularities and discrepancies between the representations that would have aided their understanding. In this section,

research examining factors that influence the way learners translate between representations is discussed and four different approaches illustrated. The first is an analytic approach specifying the nature of translation activities between representations. The second is a qualitative approach examining one subject's understanding of the mapping between representations. A third approach is to build a computational model of how an expert uses MERs. The final approaches are quantitative: one is an account of individual differences in translating between representations in one learning environment and the other uses quantitative measures to look at the influence of different combinations of representations on learners' translation behaviour. A further important issue for understanding translation between representations is how this is affected by different learner characteristics. Finally, a conceptual problem in approaches to studying translation will be considered.

3.4.1 Studying translation

Janvier (1987) provides a description of the nature of the translations between some common representations (Figure 7). He also indicates that translations between two representations are commonly achieved via a third (for example, formulae through tables to graphs). Interestingly, this may be changing with the advent of computer tools for manipulating representations - whether this change is beneficial or not is yet to be resolved. Using this table as an analytic tool, he also argues that when teaching translation between representations, these processes should be considered as complementary pairs (*e.g.* the interpretation of graphs as situations and verbal descriptions and the complement of *sketching* graphs from verbal descriptions). However, he does not provide detailed process accounts of these translation activities. Further empirical work is required to determine exactly how learners perform these operations. This is illustrated for two of these cells by the next approach.



From \ To	Situations Verbal Descriptions	Tables	Graphs	Formulae
Situations Verbal Descriptions		Measuring	Sketching	Modelling
Tables	Reading		Plotting 	Fitting
Graphs	Interpretation	Reading Off		
Formulae	Parameter Recognition	Computing 	Sketching	

Figure 7. Janvier's model of translation process between different representations (adapted from Janvier, 1987)

Schoenfeld, Smith & Arcavi (1993) examined one student's understanding of function using the Grapher environment. Using micro-genetic analysis, they describe in detail the mappings between the algebraic and graphical representation in this domain. Working with one student over a number of sessions, they showed how a student could appear to have mastered fundamental components of a domain both in terms of algebra and in terms of graphs. However, as some of the connections between these modes of representation were missing, her behaviour with the representations was often misguided. For example, she could generate the slope-intercept equation for a line, yet not realise that the x value in ' $y = x + 8$ ' would give the y value. Schoenfeld *et al.*'s analysis is most useful in that it reveals the complexity of the mappings that can exist between representations.

Tabachneck-Schijf, Leonardo & Simon (1997) describe a computational model designed to simulate an expert's use of MERs. The authors are committed to the Mind's eye hypothesis that a picture will be represented in long term memory in essentially the same form as it is perceived in short term memory. This forms the basis of CaMeRa, which consists of a pictorial external display used for reasoning and input to short term memory. Short term memory and long term memories are split into pictorial and verbal elements and network representation are used for pictorial information and propositional lists for verbal elements. Knowledge is organised so that associations between modalities are permitted but modification is only possible within modality. It also has basic semantic information about the domain in which it operates, in this case, economics. CaMeRa represents an attempt to model how important visual reasoning is to an expert and how it is tied into verbal reasoning. It also emphasis how important domain knowledge is in guiding the construction and interpretation of representations. It is not yet a model of how people learn to use and learn with MERs, although the authors point to that development.

Schwartz & Dreyfus (1993) used quantitative process measures to examine how individuals integrated information between different representations designed to teach the concept of function. They used the TRM microworld that allows users to switch between algebraic, tabular and graphical modes. They defined two measures of students' performance with the software, a *convergence index* and a *passage index*. The former describes the efficiency with which a learner uses available information to progress towards a solution. If learners progress towards a right answer quickly, then they will have a high convergence index and will be assumed to have correctly interpreted the information at each stage in the solution. The passage index describes the extent to which a student keeps track of the available information when switching between representations. Thus, a student might be described as 'Pg = (4 2+ 2-)', which states that they switched representation four times, twice transferring all the available information successfully, and twice not. Using these measures, they describe four prototypical students who differed in the success of their problem solving. For example, a student with high passage and convergence indexes was shown to be able to use the presented information successfully and keep track of it through the different representations. Another student who did not switch between representations 'Pg = (0 0+ 0-)' converged quickly on a solution through knowledge of algebraic representations alone. In contrast, less successful students had much lower convergence indices and did not pass information between representations successfully. Schwartz & Dreyfus conclude that such measures of representation use will provide useful insights into the design and use of learning environments.

Ainsworth, Wood & Bibby (1996) used similar process measures to Schwartz and Dreyfus to explore how different combinations of representations influenced learners' abilities to translate between representations. The learning environment (CENTS. see Figure 8) is designed to teach computational estimation and uses MERs to focus on the complex problem of understanding the relation between the estimate and the right answer (referred to as judgement of estimation accuracy).

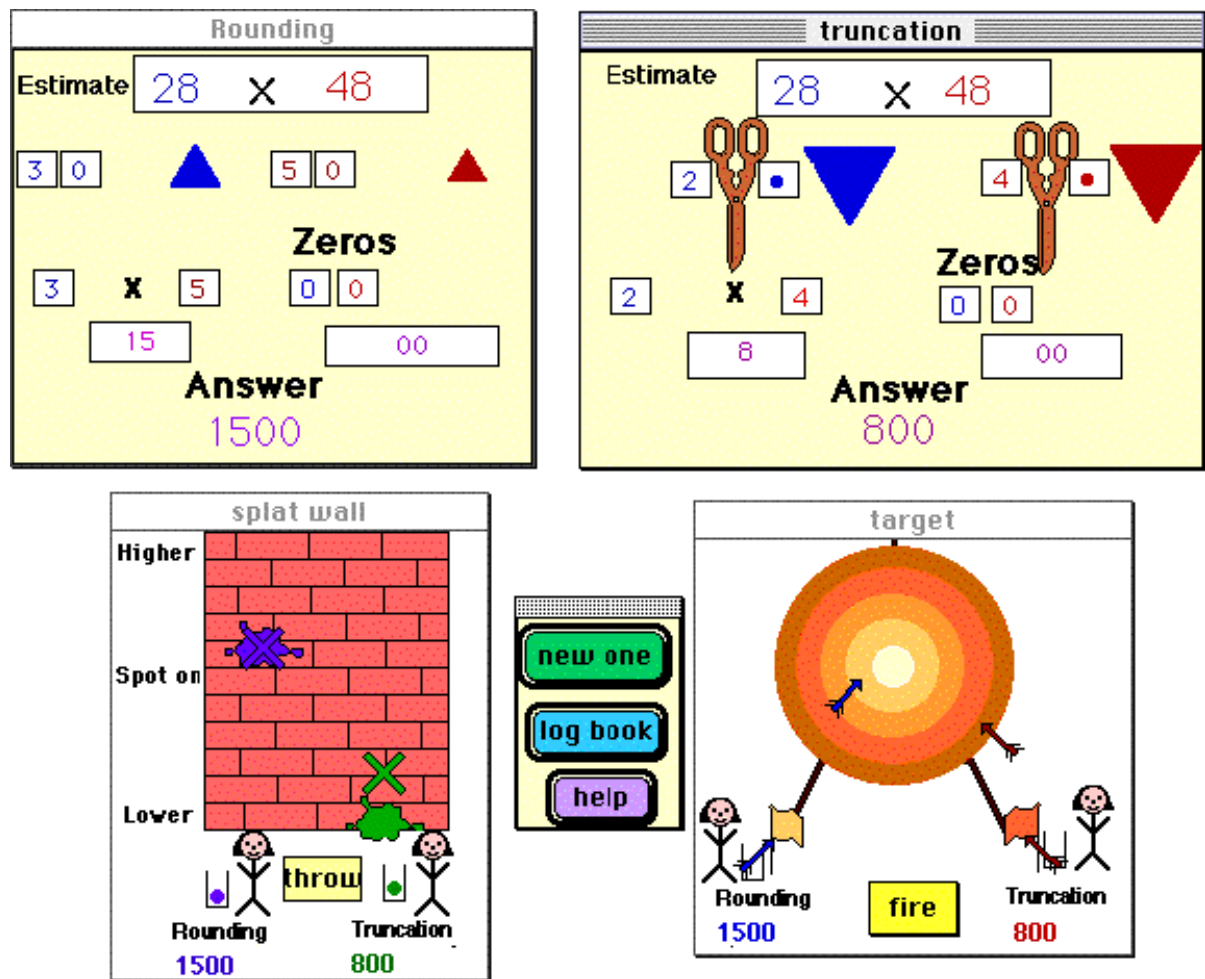


Figure 8. CENTS displaying two pictorial representations

Two partially redundant representations were used to display the direction and magnitude of the estimation accuracy. These were either two pictorial representations, two mathematical representations or one pictorial representation and one mathematical representation combined to give a mixed system. Children in all experimental groups learnt how to estimate but only children given pictorial and mathematical representations improved at judging the accuracy of their estimates. Children in the mixed condition learnt how to estimate without understanding how the estimates they produced related to the exact answer. Analysis of the process measures showed that children in the mixed condition demonstrated poorer understanding estimation accuracy in their use of each representation during the intervention. Yet, each representation in the mixed condition was also present in either the pictorial and mathematical conditions where it was used successfully. Hence, poorer performance with mixed representations can be seen to lie in the combinations of representations, rather than the in the individual representations *per se*.

To examine the influence of learning to translate between representations, the similarity of learners' behaviour over the two representations was correlated to give a

measure of how well learners understood the relation between representations (representational co-ordination). 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. Further studies showed that these problems persisted over long periods of time and even when the representations were fully redundant (Ainsworth, Wood & Bibby, 1997).

Differences between these combinations of representations were analysed to explain why translating between a pictorial and a mathematical representation proved to be so difficult for learners in these experiments. Ainsworth (1997), it was concluded that the more the format and operators of representations differ, the harder the learner will find the task of mapping between representations. An initial set of factors that maximise difference in format and operators is given in section 4.2 when the similarity between presented representations is considered.

3.4.2 Learner's characteristics and translation

In addition to representation features, a number of learner characteristics will also influence translation. Probable candidates include: (a) familiarity with the representations, (b) familiarity with the domain, (c) a learner's age and, (d) cognitive style.

If learners are already familiar with representations presented in the multi-representational environment, then they should understand (to some degree) the format and operators of representation and the relation between the representations and the domain. Then, the final learning demand of translating across the representations should occur more rapidly. Secondly, if learners are familiar with the domain underlying the representations, then again they should learn to translate between representations more rapidly. These two arguments are based on the premises that (a) the lower the learning demands are on other parts of the task, the more attention can be focused on translation and, (b) mis-interpreting any aspect of the domain or representations could lead to difficulty in understanding the nature of translation between the representations. For both of these factors, research on novice-expert differences in physics, chess, programming, *etc.* is also relevant. Generally, it has been shown that novices tend to characterise problem representations by their surface features, not their deep structure (*e.g.* Chi, Feltovich & Glaser, 1981; Adelson, 1981). Therefore, as learners generally lack expertise either in the domain or the representations they are using, they are likely to be hampered in recognising deep structural relations between representations due to their surface dissimilarity.

A learner's age may also affect their abilities to translate between representations. Often children's performance can be seen as characteristic of novices in a domain. Nevertheless, there are likely to be developmental differences that affect use of MERs. A

number of researchers have proposed that information-processing capacity increases with age. Candidates include short term memory span (*e.g.* Case, 1985), processing speed (*e.g.* Vernon, 1987) and central computing space (Pascual-Leone, 1970). One of particular relevance is Halford's description of dimensionality which is defined as number of independent items of information that must be processed in parallel (Halford, 1993). He proposed that it is not until children reach eleven years of age that they can process four-dimensional structures. If MERs exceed this capacity then children would need to re-represent the problem for example by chunking. This suggests that often younger children would require considerable experience with the representations in order to relate them successfully (see Halford, 1993 Chapter 8, for a description of this process for subtraction with Dienes blocks and the place value representations discussed earlier).

Last, the issue of cognitive style and individual differences may well be relevant. There has been much research relating both personality and cognitive factors to learning with external representations (see section 2.1.3). There is less research into aptitude-treatment interactions and MERs. An exception is that of Oberlander, Cox, Monaghan, Stenning and Tobin (1996). They suggest that one distinguishing characteristic of people who were classified as diagrammatic reasoners was their ability to translate information across representations more successfully.

The research reviewed suggests that a number of learner characteristics will affect how easily MERs are co-ordinated. Although, this is not the only factor that can determine the effectiveness of multi-representational learning environments, the impact of these learner characteristics should be considered alongside task and representations demands discussed further in section 4.2.

3.4.3 Different routes to understanding

One final issue to be considered before ending this discussion of the process of translating between representations is to distinguish the route by which a learner relates two representations.

Kaput (1989) distinguishes two ways of acting upon representations. Syntactic actions involve manipulating the symbols of a representation guided by the syntax of the symbol scheme alone, whereas semantic actions are guided by the referents of the symbols. Similarly, a learner can come to understand the relation between representations by these two different routes (Figure 9 & 10). In the first case, referred to as semantic translation, each representation is related to objects in the domain (reference field) and it is that domain knowledge that mediates the understanding of connections between the representations. In other words, it is knowledge of the represented world which supports the co-ordination of the representations. In the second case, elements of the two representations are directly related

without reference to the domain that they denote (syntactic translation) so translation occurs purely at the level of the representing world.




Figure 9. Semantic Translation




Figure 10. Syntactic Translation

The difference between these two processes can be seen in the following concrete example. Two representations are used to describe a skater's motion in a frictionless world; the first is a table of values that gives the skater's velocity at five second intervals and the second is a velocity-time graph that describes the skater's motion. A learner could either proceed by first relating each representation to domain knowledge (semantic translation). The positive gradient in the velocity-time graph and the increasing (absolute) intervals between the values in velocity column of the table both show that the skater is accelerating. Understanding how each representation describes the skater's change in speed helps the learner understand the relation between the representations. Alternatively, in syntactic translation, one route to translation is through knowledge of how the values in the table can be plotted to give the X, Y co-ordinates of points on the graph. Thus, the relation between two representations could be understood without reference to the concept of acceleration.

How these alternative approaches to translation operate is not fully understood. We do not know what role they play in contributing to learner's developing understanding nor how they might interact. Kaput suggests that there may be a progression from semantic actions upon representations to syntactic actions, but whether this is true for semantic and syntactic translation between representations is not known. The approaches to understanding translation that were described above tend to have focused on either syntactic or semantic translation or have not distinguished between the two processes. Janvier's approach focuses on syntactic translation and the passage index of Schwartz & Dreyfus does not allow us to determine the translation route or more probably combination of routes that their subjects were using. Ainsworth's paradigm of representational co-ordination suggests some syntactic translation was occurring in her experiments but cannot be used to determine the respective contribution of syntactic and semantic translation. Of the approaches to analysing translation

presented above, only microgenetic accounts such as Schoenfeld's or careful model building such as Tabachneck-Schijf *et al*'s seem to have the potential to answer this question.

Research that examined the processes by which learners translate between representations has indicated that it is a complex task. Although, there are many studies that illustrate the difficulties faced in reconciling multiple representation, a complete process model that also accounts for representation, task and learner characteristics remains to be specified. In the absence of this specific information, software that employs MERs is still being created so the task of supporting effective learning must be considered. In the next section, the design decisions that are unique to multi-representational learning environments will be described.

4.0 DESIGN DECISIONS

In order to produce effective learning environments, designers must consider a number of different issues. They must make judgements about what to teach, who the desired users of the system are and what teaching strategies the system will employ. In addition, there are specific representational questions concerning the advantages of employing certain representations for particular users and tasks needed to be addressed in order to determine which representations are most beneficial to the particular goals of the learning environment. This is a very substantial endeavour in its own right.

Issues such as these are common to all environments, but there are a set of design dimensions that uniquely apply to multi-representational systems and it is these that are considered here. They are proposed to be: (a) the way that information is distributed in the multi-representational system; (b) the similarity of the presented representations; (c) automatic translation between representations; (d) the number of representations employed and; (e) the ordering and sequencing of representations

Existing multi-representational learning environments differ along many of these dimensions. However, it can be difficult to identify exactly how they are designed as key aspects of the systems may not be reported explicitly. Consequently, the goal of making explicit the key design decisions is aimed at serving a number of functions - it will allow classification of existing systems, provide the basis for generalising empirical findings and can lead towards a set of principles for designing effective multi-representational learning environments.

In this section, the five design decisions are briefly introduced and then existing research and learning environments that address these dimensions are reviewed in more detail.

(a) Multi-representational systems allow for flexibility in the way that information is distributed between the representations and, consequently, the redundancy of information between representations. At one extreme, each representation in the multi-representational

system could express the same information (same elements of the represented world). Here, the only difference between the representations is in their computational properties. At the other extreme, each representation could convey completely different information. Multi-representational systems can also be partially redundant, so that some of the information is constant across (some of) the representations. Accordingly, one important design decision is the redundancy of information between representations.

(b) When a learning environment employs MERs, then the similarity of these representations can be very different. For example, multi-media systems can display pictures, text, sound, equations, and graphs simultaneously. Given the research reviewed above which showed the difficulty learners find in translating between representations, designers should consider how different the representations in the learning environment will appear to the users of the systems.

(c) With the advent of computer technology, it is now possible to automatically link representations in a way that was not possible with pen and paper techniques. So, the third issue that should be considered is whether to provide automatic translation between representations. Commonly, automatic translation is implemented such that a learner would act in one representation and see the results of these actions in another. Other systems require users to translate the information between the representations themselves. The research reviewed above described the considerable difficulty that learners have in translating between representations. However, it does not necessarily follow that we should provide this translation for users. It may be possible to over-automate and so deny learners the opportunity to construct knowledge of how to translate between representations.

(d) The most visible decision about the design of a multi-representational learning environment is how many representations to employ. By definition, a multi-representational environment uses at least two representations, but many systems use more than that. A related issue is how many representations to use simultaneously? Some learning environments display only a subset of their available representations at any one time. At one end of the continuum, some systems provide learners with a maximum of one representation at a time and, at the other extreme, systems such as the Visual Calculator (Fox, 1988) present five different representations simultaneously.

(e) The final design decision to be taken about multi-representational systems is concerned with the ordering and sequencing of representations. If not all of the representations in the system are presented simultaneously, a number of further issues arise. The first issue is the order in which the representations should be presented. When a sequence has been determined, then further complexity is provided by the necessity of deciding at what point to add a new representation or switch between the representations. Additionally, the designer needs to consider how many of these decisions are under system control or whether learners can make some or all of these decisions for themselves.

4.1 Distributing information between Representations

By providing MERs, a designer can choose how to distribute information between the representations. If a learning environment uses only one representation then all the dimensions of information that are presented is given by this single representation. However, if multiple representations are used, then the same number of dimensions of information can be presented in many different ways. In this section, the design space will first be set out, and then experiments that have addressed this aspect of multi-representational systems will be considered.

At one extreme, each representation could express the same information, differing only in the way that information is presented (*e.g.* a histogram or a pie chart). In this case, the multi-representational system is fully redundant as each representation refers to the same elements in the represented world. At the other extreme, each representation could convey completely different information. In this case, there is no redundancy in the system. Finally, MERs can also be partially redundant, so that some information is constant across the representations. Of course, given the different computational properties of representations, it may be impossible for a learner to derive the same information from two different representations even if theoretically the information in each is identical. However, it is to consider this aspect of representation design, as representational systems which distribute information in varying ways can serve quite different functions.

Consider the situation where the goal is to present four dimensions of information about a population such as the number in a population of a certain age, gender, marital status and nationality. Appendix One includes eight different representations of this information - each presents either one, two, three or four dimensions of information. If only one representation is used then all this information must be presented together. However, the use of MERs presents many possibilities. At one extreme, all of four dimensions of information can be presented in each representation. For example, both a table and an X, Y scatter graph could be used (see figures A1 and A2 in Appendix One). Theoretically, there is no maximum to the number of representations in the system as each additional representational will still refer to all the same elements of the represented world. Full redundancy is likely to be used when the designer wants the learner to benefit from the different computational properties of the representations and when the information itself is not too complex.

Rep A Rep B




Figure 11. Two representations that both display the same four dimensions of information*

When there is no redundancy in the multi-representational system, there are four alternative ways that four dimensions of information can be distributed between the representations (see Figure 12). At one extreme, each representation could display just one of the dimensions. Alternatively, one representation could provide the majority of the information with a second representing only one dimension of information. For example, age, gender and marital status in one representation with nationality displayed separately (*e.g.* figures A3 and A8). The other alternatives are two representations with two dimensions of information each and finally one representation with two dimensions and two further representations with one dimension each.

Distributing information in these ways allow a designer to simplify each representation. Representations that present too much information are likely to be complicated to interpret. The particular choice of representational system will depend on the inferences that the representations are designed to support. If a learner needs to know about the relation between dimensions then fewer representations with more dimensions will be required. However, if this is not necessary, then representations with fewer dimensions are often easier to interpret. For example, consider the problem of determining the number of Italians in the population with either Figure A2 or A7.

* The numbers refer to the amount of information expressed by each representation. Size of circles is proportional to this number.




Figure 12. The four possible ways that four dimensions of information can be displayed if no redundancy is permitted in the representational system

There are many ways that partial redundancy can be achieved in a multi-representational system. One important contrast is when a second representation displays a subset of the information that is presented in the first representation (Figure 13 or consider concrete examples such as Figure A1 with Figure A8, or Figure A2 with Figure A4) versus the case when both representations contain some unique and some overlapping information (Figure 14; *e.g.* Figures A3 (age, marital status and nationality) and A6 (gender & marital status)).




Figure 13. Two partially redundant systems where the second representation provides no new information. Example 1 is referred to as minimum subset and example 2 as maximum subset

Another contrast between different partially redundant systems is the degree of overlap between the representations. On the LHS of both Figure 13 and Figure 14, there is only one dimension of information that is presented in both of the representations (in the concrete examples suggested, nationality in Figure 13 and marital status in Figure 14). On the right hand side of the figures, both representations present the majority of the information and

have only one dimension that is unique (*e.g.* RHS of Figure 13 could be seen in Figure A1 and A3 and RHS of Figure 14 is illustrated by a combination of Figures A3 and A4).




Figure 14. Two partially redundant systems where both representations provide some unique and some overlapping information. Example 1 is referred to as minimum overlap and example 2 as maximum overlap

It can be seen that there are many possible ways to distribute information in a multi-representational learning environment. Existing learning environments use the complete spectrum of full, none and partial redundancy between representations. For example, MoLe (Oliver & O'Shea, 1996) uses one representation to express the relation between different modal worlds, and another to illustrate each world's content. In this case, there is no redundancy between the two representations. This choice is often made when a single representation would be insufficient to carry all the information about the domain or would be too complicated for people to interpret if it did so. Alternatively, each representation in a learning environment can display the same information (*e.g.* Block's World Thompson, 1992) where often the goal is to teach the relation between these representations. Distinguishing the types of partial redundancy used in learning environments is less easy. One system that allows different levels of redundancy is CENTS. (Ainsworth *et al*, 1996). Experiments with this system have been contrasted both full and no redundancy combinations. Other experiments have used a partially redundant system where the first representation presents both the direction and magnitude of an estimate and the second presents a subset of this information displaying only the magnitude of estimate.

Designers may have no choice concerning the degree of redundancy between representations. For example, if the goal of the system is to teach the relation between representations then full redundancy is commonly required. However, when presenting complex information, it is often possible to alter the redundancy between representations. Correspondingly, this prompts the question of whether there is a level of redundancy which best supports learning (for a particular task and a particular type of learner). One possibility is that it is easier to learn complex ideas when each part is represented separately. Alternatively, it may be harder to learn when there is no redundancy as the relation between representations (and therefore elements of the represented world) may well be less obvious.

One research paradigm that has been used to examine how redundancy influences reasoning with external representations is Sweller's cognitive load approach. He argues that as working memory is limited, material presented to learners should be structured so as to reduce cognitive load. From this position, he explores the effects of providing (informationally) redundant text with diagrams. Kalyuga, Chandler & Sweller (1998) report a number of studies that show that less experienced learners benefit from redundant text but that with more experience the same text interfered with performance with diagram. These results are interpreted as showing that learners who can benefit from the diagram in isolation do not need text and so eliminating it reduces cognitive load. This suggests that redundancy should be reduced as expertise grows. However, these experiments only explore one particular use of multiple representations - one where the text is used to constrain interpretation of an unfamiliar representation. This conclusion may not hold for other uses of MERs.

Another study that addressed how redundancy influences learning was that of Ainsworth *et al* (1997a) who drew a different conclusion. Using CENTS, two classes of representational system were created. In no redundancy situations, each representation expressed a different dimension of information. Thus, one representation was used to display direction (either higher or lower than the exact answer) and one to express magnitude (close to far, either continuously or categorically). Fully redundant MERs expressed exactly the same information in two representations - each one is (theoretically) derivable from the other. For example, both representations expressed direction and (continuous) magnitude. These two levels of redundancy were implemented as both pictorial representations and a combination of pictorial and mathematical representations. This design allowed predictions to be tested concerning how redundancy between representations affects the process of translating between representations.

In contrast to experiments with partially redundant representations in CENTS (described in section 3.4), learners with mixed representations **did** improve at translating between representations over time (although they were still significantly poorer than those given pictorial representations). However, when examining domain performance, students given representational systems without redundancy were shown to understand aspects of estimation accuracy faster than those given fully redundant representations. This experiment provides tentative evidence that initial acquisition of concepts may be facilitated when each representation expresses a different aspect of the situation, so limiting redundancy. In contrast, understanding the relation between representations is favoured by increasing redundancy. Further research is needed to clarify and extend this conclusion.

The experiments with CENTS have implications for the different types of partially redundant systems presented above. In cases where a second representation consists of a subset of the information, then ignoring this representation should not produce catastrophic

effects on learning (Ainsworth *et al*, 1997a). However, when two representations present some unique information, they must both be understood by the learners if they are to fully understand the learning situation. In the case where the second representation presents only a small subset of information (minimum subset) then attention will be drawn to that dimension of information presented by both representations. This seems less likely in cases of minimum overlap as commonly both representations will display a complex amount of information and the shared dimension may be less immediately visible. If the system goal requires learner to translate between MERs, then overlap between them may well help learners co-ordinate the representations. The experiments with CENTS reviewed above suggest that the greater the overlap between representations, the more effective this process.

The ways that distributing information between representations influences how effectively the different functions of multi-representational learning environments is considered further in section 5.

4.2 Similarity between Representations

When a learning environment presents information in MERs, these representations can differ from each other in two distinct ways. The representations can refer to different elements in the representing world (section 4.1), but commonly they also differ in the representing world level where it is the presentation of that information that varies between representation.

Research reviewed above listed the substantial difficulties that students face in learning how to use external representations and in particular the complexity of integrating information from more than one representation (*e.g.* Tabachneck *et al*, 1994). Ainsworth (1997) provided evidence that this learning demand increases as the difference between the representations increases and showed that when the difference becomes too great, it can inhibit effective learning. This problem provides the designer of multi-representational software with a difficult task - how to exploit the properties of different representations without over loading the learner with impossible costs in translating between the representations.

In order to address this question, two research questions need to be explored. Firstly, the specific role that alternative representations play in supporting learning needs to be understood. Secondly, designers should consider what factors affect learner's abilities to translate between representations in order to minimise unnecessary complexity. The first research question has been addressed extensively in the literature on learning and problem solving with external representations (see section 2.1.3) and systematic answers have begun to emerge. However, the second issue has been less actively researched and a framework for addressing this question will be presented in this section.

Research reviewed in sections 3.3 and 3.4, showed how complex mapping between different representations can be for learners. For the most part, this research has considered

representations with very complicated mapping rules such as equations, tables and graphs (e.g. Yerushlami, 1991; Schoenfeld *et al*, 1993; Schwartz & Dreyfus, 1993). Given the sophisticated knowledge needed to relate these representations (see Schoenfeld *et al*, 1993 for a detailed account), it is not surprising that experiments with learning environments employing these representations have shown that integrating these representations is a difficult task for many learners.

Research on analogical reasoning (e.g. Gick & Holyoak, 1980; Gentner & Toupon, 1986) has examined factors that influence people's abilities to recognise similarities between problems. It has consistently been shown that people find it difficult to recognise the similarity between problems as they become misled by surface dissimilarity and fail to recognise commonality in the deep structure. This has particular relevance for learning as opposed to problem solving as novices tend to categorise problems by their surface features (e.g. Chi *et al*, 1981). To provide a definitive statement of the factors that affect learners coming to understand the relation between representations awaits further research and possibly an integrative taxonomy of representations. However, based upon current evidence, it is possible to state that these factors will involve both learner and representation variables for any given task. The issue of learner characteristics has been considered in section 3.4, and the issue of how best to characterise tasks is important, but outside the scope of this paper. In this section, representation factors will be considered.

Representations consist of information (the represented world) and symbols to display it (the representing world). Hence, logically, there are four classes of multi-representational system. The first situation is where two representations describe exactly the same information but do so in different symbol systems. For example, the distance travelled by a body in a frictionless world is given by the equation $s = ut + 1/2at^2$ or by finding the area under a velocity-time graph or reading the maxima of a distance-time graph. These allow someone who is familiar with the representations to determine distance travelled but have very different ways of expressing this information (same elements in the represented world, different elements in the representing world). Whereas, the velocity time-graphs of two different bodies are identical at the symbol level but express different knowledge (different elements in the represented world but same elements in the representing world). Thirdly, a velocity-time graph of body one and a distance-time graph of body two; express different information in alternative ways (different representing and represented world). Finally, it is more trivially possible for two representations to express the same information in the same way (same representing and represented world).

In the next two sections of this paper, aspects of representations that can be related to (first) the representing world and (second) the represented world will be described. The dimensions have been extracted from the research literature and should not be considered

exhaustive. They are intended to form the basis of more principled research into these factors.

4.2.1 Representing world

It has been proposed that the more the formats of representations and the operators that act upon them differ, the more difficult it will be for learners to recognise their similarity. The following set of dimensions is suggested as likely to maximise these differences. The first six items refer to differences in the format of representations (and hence their operators) and the second four items more specifically to differences in operators as the format of these representations need not necessarily differ.

- the modality of the representations - propositional v graphical
- the levels of abstraction (*e.g.* concrete to symbolic representations)
- the type of representation (*e.g.* histogram, equation, table, line-graph)
- the specificity of representations
- whether representations are static or dynamic
- differences in labelling and symbols on the representations
- alternative uses of representations, *e.g.* display v action
- the interface to the representations
- self-constructed & selected representations versus pre-determined representations
- whether the representations encourage different strategies

The most traditional distinction between representations is that of modality where graphical/diagrammatic representations are contrasted with sentential/propositional. The important difference is that graphical representations explicitly preserve geometric and topological information. This distinction has formed the basis of research into the properties of diagrammatic representations and their advantages in solving certain types of problem (*e.g.* Larkin & Simon, 1987; Stenning & Oberlander, 1995). It has also been central to the design of many learning environments (*e.g.* HyperProof, Barwise & Etchmendy, 1995; MoLe, Oliver & O'Shea, 1996).

The degree of abstraction of a representation has been considered by many researchers. One classic distinction is that of Bruner's (*e.g.* Bruner, 1966) who described three different modes in which knowledge is expressed - enactive, iconic and symbolic. Enactive representations are physically expressed, iconic representations are pictorial in nature and bear a one to one correspondence with the objects they represent and symbolic representations have an arbitrary, non-perceptual relationship to the object they depict. Purchase (1998) adapts Bruner's scheme by dividing the iconic category in two: concrete-iconic which has a direct perceptual relationship to the object (*e.g.* a photograph) and

abstract-iconic which has a related but non-direct relation (*e.g.* a road-sign warning of falling rocks). Another similar scheme based on increasing abstraction of representations is Fieldman's (1993). She proposed six levels of fidelity: depictive/pictorial representations are the most realistic and include 3d models and colour photographs; schematic representations include maps and caricatures; iconic representations include international signs and hieroglyphics; structural/functional representations retain some fidelity to their referents and can be seen in flow charts, blueprints and graphs; symbolic representations include logos and symbols such as skull and cross-bones and finally; arbitrary representations have no perceptual relations to the objects depicted and include tables and text. It can be seen by the variety of classifications that exist for this dimension that the granularity of these distinctions is to some extent arbitrary. Blackwell & Engelhard (1998) list eight and they do not include any of the ones proposed here, Its persistence in classifications suggest that many researchers find this distinction a crucial way of distinguishing between representations.

Many different ways of categorising representations into different *types* have been proposed (*e.g.* Lohse, Biolsi, Walker & Rueler, 1994; Lesh *et al.*, 1987; Kaput, 1987; Cox, 1996). For example, Lohse *et al.* identified eleven major clusters: graphs, numerical and graphical tables, time charts, cartograms, icons, pictures, networks, structure diagrams, process diagrams and map clusters. Taxonomies such as these have been created by a variety of methods (*e.g.* intuition, analysis of domain properties and card sort techniques with subjects) and by researchers with a variety of backgrounds such as cognitive science, education, HCI, graphical design and psychology to name but a few. Although there is some overlap between the taxonomies, no one classification is universally accepted. They differ in the domains addressed, the granularity with which representations are described and the task for which they were created. Consequently, it is problematic to state that representations that are different in type will be harder for learners to reconcile, as different type means contrary things to these different researchers and importantly to different learners. Nevertheless, it does not seem a worthless distinction to make. If learners consider that two representations are of alternative types they may well manipulate and interpret them differently. Accordingly, if a well-established taxonomy exists for the domain/representations that the learning environment addresses, it can be used heuristically to consider how learners will interact with representations.

Stenning & Oberlander (1995) identify *specificity* as a fundamental property of a representation that has direct ramifications for processing efficiency. Specificity is the demand by a system of representation that information in some class be specified in any interpretable representation. The specificity of a representation determines the extent to which the representation permits expression of abstraction. Stenning and Oberlander propose that there are three main classes of representation: Minimal, Limited and Unlimited Abstraction Representational Systems (MARS, LARS, and UARS, in increasing order of

expressiveness). It is proposed that the class that each representation belongs to will predict their cognitive computational properties, with a LARS being more computationally effective than a UARS as these systems are syntactically constrained and limit the number of cases that must be computed over. This analysis again predicts that learners will interpret and act upon different types of representations in different ways. Hence, this is likely to increase the chances that similarities between representations will be missed by learners.

One way that designers can signal relations between representations is to be constant in their use of labels and codes across representations. This problem was highlighted by DuFour-Janvier, Bends & Belanger (1987). They describe children who were happy to get three different answers to the problem $3152 - 128$ when solved using an abacus, vertical calculation and horizontal calculation. These researchers suggest each representation became a different problem to the learner and so it was not surprising that they produced different answers. They propose that children only have a tendency to recognise that two representations concern the same problem if they contain the same numbers. Thus, the numbers on the representations acted as labels to help children translate between the representations. Good design can be seen when labels on buttons are consistent when their function is the same (*e.g.* just one from quit, exit, bye and leave) In road atlases, route finders label motorways and 'A' roads in blue and red respectively. The finer-grained maps also keep to this labelling. Inconsistencies in labelling will make it harder for learners to relate representations.

The introduction of information technology into the classroom has brought a new type of representation to learning situations - dynamic representations. These include animations which have been defined as a series of rapidly changing static displays giving the illusion of temporal and spatial movement (Scaife & Rogers, 1996). For example, a typical educational application of an animation is of blood flowing around the body within a biological CD-ROM (Jones, 1998). Dynamic graphs are also becoming more common. Experiments can be run either in simulation such as microworlds designed to explore Newtonian motion or with datalogging equipment and graphs updated as the experiments progress. These sorts of static and dynamic representations require different operators to interpret them and have different formats. This can affect what students learn with these types of representations. Jones (1998) showed that the sorts of errors made by students interpreting blood flow round the body were affected by whether they saw an animation or static diagrams. Dynamic and static representations support different inferences and so when presented with both types of representations, learners may well have difficulties in reconciling them.

A related issue is the way that information technology can alter the role of a representation from being that of display to action. This difference is due not to absolute properties of the representations, but to features that evince different patterns of use. Display representations are not intended to be acted upon by users, except to be built initially. Action

notations support a variety of transformations and actions. For example, transforming equations, substituting values for variables and extending tables are all examples of actions upon representations. Traditionally representations such as line-graphs and histograms were display representations as they were fairly laborious and time consuming to construct with pen and paper. However, computer technology such as that employed in graphical calculators or in learning environments such as FunctionProbe (Confrey, 1992) now allows users to act upon and manipulate these types of representations. This offers new and exciting educational possibilities for these types of representations. However, it also suggests that students will have to learn how to manipulate these new features of representations. As the processes involved in working with display and action notations are different, when presented together within a multi-representational environment, it may be difficult for learners to translate information between them.

The way that representations are acted upon differs not only in terms of the display/action properties of the representations described above, but also in terms of the interface to these representations. A common design decision is that representations that are propositional tend to be acted upon via the keyboard, but those that are diagrammatic are accessed using direct manipulation devices. Recent research has demonstrated that the choice of interface can influence what users learn. For example, Svendsen (1991) found that direct manipulation interfaces resulted in poorer performance than command lines interfaces for solving Tower of Hanoi problems. Consequently, some researchers are now arguing for a move from direct manipulation interfaces in educational technology (*e.g.* Gilmore, 1996) or for more attention to be paid to the way that actions on representations are supported by educational technology (*e.g.* Churchill & Ainsworth, 1995). This is true for MERs where the use of more than one interface style may serve to increase perceived differences between representations.

Learning environments can use representations in three different ways. Representations can be constructed by the learner with free choice about how and when to use them. Alternatively, pre-fabricated representations may be presented to the learner as the system chooses. Finally, partially constructed representations might be available where, for example, a learner is given a table but must fill in values for themselves. The processes involved when constructing representations are very different to those where the learners task is to interpret presented representations. In the first instance, learners must interpret the problem, choose the representations, construct the representation and interpret the representation before responding (Cox, 1996). In the second case, the learners must interpret the problem and the representations before responding. The interpretation of the representations should be much easier if the representations were self-constructed as learners should be familiar with the rules governing the representations. However, this can also introduce an additional source of error if representations were constructed mistakenly. It can

be seen that the processes involved in constructing representations are very different to those involved for interpreting them. This claim has been made for a number of domains. For example, Piroli & Anderson (1984) showed that teaching students to interpret recursive LISP code did not seem to help them write it. This is given a theoretical basis in architectures such as ACT* which emphasises the use specificity of production rules (*e.g.* Anderson, 1989). Once again, the purpose of this paper is not to argue which is the better use of representations (or even the more subtle question, which is the better use representations for which problem and which user) but to identify factors which affect how students reconcile MERs.

One of the reasons for using MERs is that by doing so learners can be encouraged to use different strategies. As described above (section 2.1.3), a number of researchers have shown that different representations promote different strategies (*e.g.* Tabachneck *et al*, 1994; Cox, 1996). This effect has also been observed in children (*e.g.* Watson, Campbell & Collis, 1993). There seems little doubt that different representations can encourage alternative strategies and that often this can be advantageous as by switching between representations learner can compensate for weaknesses in the strategy or representation. However, if learners are attempting to relate different representations then this may provide a source of difficulty. Ainsworth (1997) hypothesised that one of the reasons why learners did not integrate information presented in pictorial and mathematical representations is that the pictorial representation encouraged the development of a perceptual strategy and the mathematical one encouraged learners to generate a rule based upon symbol manipulation.

4.2.2 Represented world

The second way that external representations can differ is if they express different information (elements of the represented world). Two such ways are (a) amount of information per representation and (b) variations in the precision of presented information.

The *amount of information* has been defined as the number of dimensions of the represented world that a representation encodes (Palmer (1978) refers to this as type of information). For example, in CENTS some representations display either the direction or the magnitude of the estimates (one dimension) and some both direction and magnitude (two dimensions). Appendix One provides examples which described a hypothetical population with four dimensions of information (age, gender, marital status and nationality). In these examples, the number of dimensions of information is unambiguous. However, with more complex domains it may be harder to find a consistent way to characterise the information given by the representations.

The amount of information in a representation is particularly interesting when considering multi-representation systems because it allows for different levels of (informational) redundancy across representations. Three levels of redundancy are possible - no redundancy, varying amounts of partial redundancy and full redundancy. Ainsworth *et al*

(1997b) present tentative evidence that increasing the redundancy between representations, allows learners to translate between representations. This issue was considered further in a separate section above (4.1).

Precision refers to the grain size of a representation. If a dimension describes n relations, the higher the value of n , the higher the resolution and the smaller the grain size. For example, someone could be described as either shorter or taller (2 relations) or 5 feet 2 inches, 6 feet inch, 165cm or 1 m 65 cm (up to an infinite number of relations) *etc.* A representation that uses continuous measurement such as centimetres could be used to derive the representations which uses shorter or taller but not vice-versa. A common distinction related to precision is the distinction between quantitative and qualitative. For example, Plötzner, Spada, Stumpf & Opwis (1990) described four levels in understanding classical mechanics illustrated in increasing precision: (a) the magnitude of \underline{F} and the magnitude of \underline{a} are related; (b) if the magnitude of \underline{F} increases, then the magnitude of \underline{a} also increases; (c) if the magnitude of \underline{F} increases by some factor, then the magnitude of \underline{a} also increases by the same factor and; (d) the quotient of the magnitude of \underline{F} and the magnitude of \underline{a} is constant - $F = ma$. Representations that differ in precision may be harder for learners to reconcile.

4.2.3 Summary

These two subsections have proposed that existing research that describe differences between representations can be used to suggest factors that affect how similar two representations appear to a learner and consequently how easily they can translate between them. A large number of factors were proposed that dealt with differences in surface characteristics of representations and it was also suggested that as novices are particularly affected by differences in surface characteristics. A fewer number of factors were associated with differences between representations at a deep level. In addition, these factors should be considered in relation to the learner characteristics that affect translation described in section 3.4. These factors could be used in two ways. Firstly, they may be used as heuristics by developers considering the design of multi-representational learning environments and by teachers considering how to support such learning (see Ainsworth, Bibby & Wood, 1997b for consideration of the teacher's role). Secondly, they could be used to guide more systematic research into the factors that affect learners' perceptions of multi-representational environments and hence their influence on ease of translation.

4.3 Automatic Translation

One question facing designers of learning environments is whether to provide automatic (dynamic) linking between representations. Here, one acts in one representation and sees the results of these actions in another. Thus, it is hoped that the relation between the representations is made more explicit and hence understandable to learners than has traditionally been possible with static media. Indeed, Kaput(1989) points to the dynamic

linking of representations as one of the most important roles for new technology in mathematics learning.

The opportunities for dynamic linking provided by the advent of computer technology have been exploited by many different systems. A classic example of this is seen in the Dienes blocks microworld (Thompson, 1992). This learning environment aims to help primary aged children understand basic arithmetic by setting sums and giving feedback in one representational system while allowing them to act upon another. Its purpose and representations are very similar to the ones that Resnick & Omanson used that were described in section 2.3.1. Children might be required to find the total of $1245 + 452$ by manipulating graphical representations of Dienes blocks (where the first sum would be represented as 1 cube, 2 flats, 4 longs and 5 singles and the second as 4 flats, 5 longs and 2 singles) with continuous feedback provided by both the standard numerical display and the expanded language given above. A common multi-representational system that provides dynamic linking is the graphical calculator that presents linked algebraic expressions and graphs. Users can act upon either of these representations and the other will change in line with the manipulations

Many researchers have declared that there are substantial cognitive benefits from the dynamic linking of MERs (*e.g.* Scaife, Rogers, Aldritch & Davies, 1997; Kaput, 1987). The underlying reasoning behind the claims is that the dynamic linking of representations is alleged to reduce the (cognitive) load upon the student. It is hoped that by requiring the computer to perform the translation activities, that students are freed to concentrate on their actions upon representations and their consequences in other representations. Kaput (1992) points out that this may be particularly beneficial when the representations involved are expressing actions sequences rather than just final outcomes as previous research has shown just how difficult this task is for learners (*e.g.* Resnick & Omanson, 1987).

The difficulty that learners have in translating between representations is undeniable. However, there are reasons to hesitate about the invariable dynamically linking representations. If the aim of instruction is to encourage users to understand the mapping between representations and to translate between them, then we may be in danger of over-automating the process. This over-automation may not encourage users to actively reflect 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 to examine the goal of teaching translation between representations from the perspective provided by scaffolding and, in particular, contingency theory (Wood, Bruner & Ross, 1976). 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. Wood *et al* describe five levels of support ranging from general encouragement (level 1), through increasingly specific help to ultimately showing the learner a solution (level 5). Wood *et al* have shown that subtly adjusting the level of help given to learners depending upon their performance is much more successful than swinging between extremes of support or sticking to one level. On this view, providing automatic translation between representations is constantly giving support at level 5 and is therefore not an ideal way of teaching the relation between representations. By the same token, providing two completely independent representations leaves the learner with no specific help and would equally not be considered an effective teaching strategy.

An alternative to both of these extremes is to signal the nature of the mapping between the representations in order to support a learner's co-ordination of the representations without actually automating the process. This is often the only way to identify linkage in a system if both representations are used for display and not action. For example, COPPERS (Ainsworth, Wood & O'Malley, 1998) presents primary school children with two representations of their answers to coin problems. The first one is the canonical place value representation and the second tabular representation that is more unfamiliar to this age group. Highlighting is used to indicate how the elements of the place value representation correspond to entries in the table (see Figure 3, page 9). It is known that tabular representations are difficult for children of this age (*e.g.* Underwood & Underwood, 1987). Therefore, highlighting is intended to provide learners with some support as they come to understand an unfamiliar complex representation that offers a new perspective on the problems that they are solving. The signalling in this case is automatically generated by the system, but an alternative is to place this support under learner control. For example, given a table of (X,Y) co-ordinates, users may wish to select a row in the table and then be shown the equivalent location on a graph as in the record screen of SkaterWorld (Pheasey, O'Malley & Ding 1997). This level of support for translation between representations represents a mid-position between dynamic linking and no linking and as such may be appropriate for certain learning situations. Again, it is susceptible to the criticism that it is not contingent upon a learner's performance. For some learners, this level of support may be redundant, but for others, it may not be enough.

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. Initially, learners could work with the familiar representation to receive feedback in a less familiar representation. 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 (acting as the dynamic linking did initially), then they should be able to work independently on either representation.

One problem with this approach is that it assumes that the representations are fully redundant. Conditions for dynamic linking of representations vary with different degrees of informational redundancy. When there is no redundancy in information between the different representations, then automatic translation is, of course, impossible. In fully informational redundant systems, there is little difficulty in linking representations as each one representation can be derived from the other. Even here though there can be problems. For example, tables of X,Y co-ordinates and graphs are often stated to be informationally equivalent but in practise this may not be the case for each instance of the representations. In a table of X,Y co-ordinates, every value in the table can be seen on the graph, but this need not necessarily be true in reverse. Thus, if a user selected a row in a table, this could be indicated on the graph, but a selected point on the graph may not be presented in the table of values.

Some learning environments are designed to display partial redundancy between representations. For example, in CENTS (Ainsworth *et al* 1996, 1997a; described above), one representation can present both direction and magnitude (D & M) of error in computational estimation sums, whereas a second presents only magnitude information (M). A learner's actions on the D & M representations can be reflected in the M representation by simply ignoring the direction information. But, if the learner is working with the M representation, then extra direction information is needed to map this representation into the D & M representation. Consequently, in situations where there is a difference either between the amount of information in the representations or in the resolution of that information, then to dynamically link these representations, learners would be required to work with the representation with the most information or to provide extra information to disambiguate their intended action.

It has been argued that to design the most effective level of support for translation between representations there is a need to monitor learner's understanding of the relation between representations independently of domain knowledge. In order to provide an appropriate degree of support, it would be useful to know how difficult learners are likely to find translating between any two given representations. The issue was addressed earlier (section 4.2) where it was proposed that the more the format and operators of a representation differed then the more difficult it will be for learners to translate between these representations. In addition, a series of learner characteristics that influence this was also provided (section 3.4). However, if the learning environment is to respond to a student's growing understanding, it is necessary to monitor this understanding dynamically as the user interacts with the system. A periodic measurement of the similarity of user's behaviour on both representations (Ainsworth's representational co-ordination) or how much information is transmitted between representations (Schwartz & Dreyfus's passage index) could be included

within the student model of an ITS. This would allow the ITS to monitor and adjust how much support it is providing for translation between representations.

There is no simple solution to how best to support translation between representations in multi-representational learning environments. It has been argued that there is a need to vary the way that this support is implemented. The relationship between the degree of automatic translation and the best way to support the different purposes of MERs will be addressed in the fifth section of this paper.

4.4 Number of Representations

By definition, multi-representational learning environments employ at least two representations. The question of how many representations should be used depends upon how the MERs serve the learning objectives of the system. These different functions were described above as providing different information and processes, constraining interpretation and deeper understanding. In all cases, it is assumed that given the learning demands associated with MERs, it is wise to use the minimum number of representations that are consistent with the function of the representations. In many cases it may not be appropriate to use MERs at all, since one representation may be sufficient. For example, Ainsworth *et al* (1997a) showed that when children were given partially redundant representations, a highly effective strategy was to ignore one of the representations to concentrate upon a single useful representation. Chandler & Sweller (1992) argue that often one integrated representation is better than two (see section 4.1).

The first question is how many representations to use to maximise the learnability of complex information by distributing it over more than one representation. This remains difficult to answer as much more research is needed on this issue. Research reviewed in section 4.1, suggests that initial acquisition of complex topics may be better served by using combinations of simple representations which contain less information rather than one complex representation. But, Kalyuga *et al* who used MERs for different purposes suggested that one integrated representation may be better for novices. More research is needed on this topic. The question of the number of representations which best support computational processes can be addressed by knowledge of the difference inferences supported by each representation and the domain to be learnt. Research such as Bibby and Payne's which showed how three representations each served a different function in learning to operate a simple control panel could provide the basis for this sort of analysis. However, this task is likely to be much more complex for real world learning situations.

The second major use of MERs is to constrain interpretation. These cases are illustrated by using a concrete/familiar representation to support interpretation of an unfamiliar representation or by exploiting inherent properties of a representation to make new inferences about a second representation. As in both of these cases the aim is to

minimise learning demands, ideally this constraint should be achieved by the addition of only one extra representation.

Three ways that MERs encourage deeper understanding were described. The first use is to explicitly teach relations between representations. By definition, environments that aim to help learners translate between representations must (at least) use those number of representations. Similarly, multiple representational systems that aim to teach extension over representations also (minimally) employ that number of representations. The final example given of deeper understanding was to use MERs to encourage abstraction. Little is known about how to use MERs to encourage abstraction without providing impossible learning demands. This issue is addressed further in section 5.3.

4.5 Ordering and Sequencing Representations

Many multi-representational learning environments present only a subset of their representations at one time. At one extreme, a number of environments display only one representation at a time (*e.g.* TRM Schwartz & Dreyfus, 1993; SwitchER Cox, 1996). Others present many simultaneously such as ReMIS-CL (Cheng, 1996a) present one, two or four at a time from a total set of seven representations. In these circumstances, two main decisions must be made - in what order to present the representations and when to change the representations that are displayed. These decisions either can be made by the designer in advance and embedded in the system or can be decided more dynamically by the learner.

A continuum of approaches can be seen to the problem of deciding a sequence of representations. At one end, some researchers start with an analysis of the properties of the domain to be taught in order to identify any representational consequences. Only in the absence of any particular constraints arising from this domain analysis are more general representational factors can be considered. Alternatively, a representational perspective can be taken which favours a domain general approach.

An illustration of the first approach can be seen in Kaput (1994) who describes a system, MathsCar, which teaches introductory calculus. He argues that understanding is best supported by introducing integration before differentiation. Consequently, he proposes representations such as velocity-time graphs should be introduced before position-time graphs. At a mid point on the continuum lies approaches such as Plötzner (1995) and Spada & Plötzner (1997). They analyse domains such as one-dimensional motion in classical physics problems and argue for the importance of qualitative knowledge in solving these sorts of problems. They go further to argue that qualitative knowledge should be taught before quantitative knowledge and consequently qualitative representations should be introduced before quantitative ones. Evidence for this proposal was initially provided by the development and evaluation of a cognitive model (SEPIA) and subsequently in examining

collaborating pairs taught with different sequences of qualitative and quantitative representations.

Alternatively, more domain general approaches to sequencing MERs can be seen in the design of multi-representational learning environments. For examples, many environments introduce representations in such a way as to increase the abstraction of the representations (*e.g.* the QUADRATIC tutor). One system already described in this paper that takes this approach is that of COPPERS (Ainsworth *et al.*, 1998) which presents coin problems to children. Problems are presented first as pictures of coins whose total can be found simply by addition. Then increasingly abstract representations are displayed, for example, mixed text and pictures and then text only. Finally, problems are presented using an algebraic notation. If children request help, then problems are rephrased in increasing concrete terms reversing this order of presentation. Such approaches are generally considered to follow Bruner's view of representation, where symbolic representations replace iconic representations that in turn have superseded enactive representations. However, this treats these representations as following a simple linear sequence which Bruner did not intend (Behr, Harel, Post & Lesh, 1992). Although, this approach can be seen in many systems, its validity has rarely been evaluated. Given the difficulties in translating between representations that differ in format (described in detail above), it is almost certainly the case that the difficulties that learners have in moving between concrete and abstract representations which have different format and operators and encourage alternative strategies have been underestimated.

However, this approach does make sense when analysing the process of learning with MERs. One way to approach the problem of sequencing in the absence of domain specific information is to analyse the representations in terms of their formats and operators and in the complexity of relating the representation to the domain to be learnt. Given the increased cognitive load associated with beginning new learning tasks, then it seems reasonable to start by offering learners the least complex available representations. Commonly, this may be the most concrete/least expressive representation that the increasing abstraction route suggests. However, this need not necessarily be the case in every situation.

Even if a sequence of representations has been determined then designers are still faced with the question of when to change a representation or introduce a new one. One possible solution is to allow learners to make this choice. For example, the SwitchER system (Cox, 1996) allows users to move at will between their self-created representations. Cox argues that this can be beneficial as it can help learners to resolve impasses. However, he found evidence to suggest that switching between representations can also be symptomatic of less understanding. Another possibility is that learners should switch when they have exhausted all of the information available in the representation they are currently using. Graphs and Tracks (Trowbridge, 1989) exploits this technique to good effect. For example, help provided by the system suggests that users should switch from a velocity-time to a

distance-time graph in order to gain information about the represented object's starting position.

Alternatively, the system may take responsibility for determining when to change the representation. In this case, the task for the system is to determine when users have learnt all they can about the domain with the given representations, but not switch so soon (or so often) that the learning demands of the new representations overburden the user. Alternatively, as Resnick & Omanson (1987) observe, it is possible to introduce new representations too late. In their study of children learning to subtract using the standard written symbols, Dienes blocks were introduced to help children understand this task in a more conceptual way. These researchers were disappointed by how little children referred to the blocks and suggested that once children had reached automated performance with symbolic manipulation that it does not easily allow for application of principled knowledge. If this finding generalises to other domains, it suggests that a new representation should be introduced before learners have achieved automated performance with an existing representation.

One possible solution to this dilemma is to provide a new representation when the learner's behaviour is still flawed with respect to the domain but has converged over the current representations. This could be monitored by techniques such as representational co-ordination or the passage index of Schwartz & Dreyfus. When these measures indicate that learners have mastered relations between existing representations, the addition of a new representation could help debug misconceptions or introduce new knowledge without overburdening users. This information could also be given to students to guide their decisions if they have responsibility for selecting new representations.

5.0 DESIGNING FOR DIFFERENT FUNCTIONS OF MULTIPLE REPRESENTATIONS

In this paper, a functional taxonomy of MERs, the learning demands associated with using MERs and the design decisions that uniquely apply to multi-representational learning environments, have been described. The task that remains is to consider how these separate factors can be integrated in order to propose design principles for supporting learning with MERs. The basis of effective design is taken to be the functions that MERs can serve. In this final section, an idealised multi-representational learning environment will be proposed for each of these functions aimed at minimising the learning demands of MERs. Due to the limitations on the current state of our knowledge concerning learning with MERs, these designs are proposed as the basis for further empirical work, not as well developed principles. Furthermore, it is unlikely that a multi-representational system will serve only one of these functions. Allowing for multiple functions will make the design decisions more fraught and may well suggest that aspects of the system should change over time. This is considered further after each functional design is discussed.

5.1 Designing for Different Information and Processes

Three distinct functions were proposed for this use of MERs: that MERs are used to convey different information, to support new inferences by providing partially redundant representations and to support different processes. As is illustrated in the following sections, the predicted designs for each of these functions are relatively similar. They are based on assuming that if each representation plays a relatively independent role, learners should not be required to learn to translate between representations.

5.1.1 Designing MERs to convey different information

The first reason to include MERs in a learning environment is to distribute information over the representations. This use of MERs is common when one representation is insufficient to carry all the required information or would be very complicated to interpret if it did so.




Figure 15. Using MERs to convey (completely) different information

Figure 15 shows an abstract illustration of a learning environment that supports this form of MERs. Each representation in the system describes a different element in the represented world. Note that there is no translation between the representations. The distance between the representation and the domain is intended to indicate the cognitive effort required to successfully use a representation.

Table 1 Design aspects of using MERs to convey (completely) different information

	Using MERs to convey different information
Redundancy	No redundancy between representations
Similarity of Representations	Maximise similarity
Automatic Translation	No automatic translation
Number of Representations	Between two and as many reps as there are dimensions of information
Order of Representations	Limit co-presence. Order determined by the task demands or learner characteristics

This particular example assumes an idealised case of no overlapping information. Hence, two of the design decisions for this function of MERs follow logically from its definition. There is no redundancy between the information presented in each of the representations. Accordingly, it is impossible for the learning environment to automatically translate between the representations as they share no information in common. The minimum number of representations in a multi-representational system is two. The theoretical maximum number of representations in a system with no redundancy is equal to the number of dimensions of information if one assumes that each representation displays only a single dimension of information. However, it is likely that a system will compromise between these two extremes by providing representations that have multiple dimensions of information. Exactly how many representations used in a system will be a compromise between balancing learning tasks and the inferences the representations are required to support. Limiting the number of representations should reduce the learning demands associated with each additional representations. But, representations that include too much information can be difficult to interpret. If a learner needs to reason about the relations between dimensions of information (*e.g.* see A1, A3, *etc.* in Appendix One) then less representations with more dimensions of information will be required.

The proposals to limit co-presence of representations and maximise the similarity of representations are aimed at minimising the need for and the cost of translation between representations. For this function of MERs, it is proposed that it is sufficient for learners to understand the format and operators of each representation and the relation between the representation and the domain. Accordingly, as translation is both difficult and unnecessary, it should be discouraged. Najjar (1998) suggests that presenting text and graphics simultaneously rather than sequentially provides learners with more opportunities to build links between the representations. However, Ainsworth *et al* (1996, 1997a) showed that when one representation was sufficient to learn the desired aspects of a domain, presenting it

alongside a second representation could interfere with successful learning and that this was due to the cognitive demands of translating between representations. Therefore, it is argued that learners are more likely to try to co-ordinate representations that are co-present and so, for this use of MERs, that sequenced representations should be preferred. A similar argument is made for similarity between representations. As with all decisions about similarity, there often may be specific reasons to include representations with certain computational properties. However, in the absence of any definite objectives, then representational systems that maximise similarity (refer to section 4.2 for a list of these) should reduce the learning demands of co-ordinating representations if learner's attempt this task.

5.1.2 Designing MERs to support new inferences by providing partially redundant representations

MERs can be used to support new inferences when the information that is partially redundant over two (or more) representations is integrated by a learner. In the discussion in section 4.1, two types of partial redundancy were determined- redundancy by subset and redundancy by overlap. For this use of MERs, each representation contributes some novel information and hence redundancy by overlap is required. The example provided in section 4.1 was navigating by using both a London Underground map and a street map. The learner is required to integrate the information provided by each representation as each contributes some unique information. But, as this can occur at the domain level there is no need to translate between the representations themselves.




Figure 16. Using MERs to support new inferences by providing partially redundant representations

Table 2. Design aspects of using MERs to support new inferences by providing partially redundant representations

	Using MERs to support new inferences by providing partially redundant representations
Redundancy	Partial overlap between representations
Similarity between Reps	Similarity may be determined by computational properties, but otherwise aim to maximise similarity
Automatic Translation	Partial or no translation
Number of Representations	minimum two, maximum unbounded
Ordering of Representations	Co-presence of the partially redundant reps

The similarity of the representations will be primarily determined by the best way to display the information to the learner. Although, both the information provided by the street map and Underground map could be given in lists of propositions, the existing formats are more effective. Consequently, in many cases combinations of representations will differ in their perceived similarity. Theoretically, this should not matter as it is proposed that integration need not occur on the level of the representing world. In practice, it may do so if learners are using their understanding of the syntactic relation between representations to understand how to integrate the information.

Automatic translation is only possible on the dimensions of information that are shared by two representations. For example, with a computerised version of the street map and Underground map, highlighting an Underground station on one map could cause it to be identified on the second map. However, translating information about streets or underground lines would not be possible as this information occurs in only one of the representations. Accordingly, automatic translation will be at most partial and may not be possible at all between some representations and dimensions of information. Where feasible, it seems likely that providing automatic translation will aid learners in their attempts to integrate the information and so is to be desired.

The number of representations required for this use of representations must balance the competing demands of interpreting complex information. For learners to integrate the crucial dimensions of information between the representations, they must be able to clearly identify the redundancy between the representations with the overlapping dimension(s). If too many representations are used then learners may not know which representations to focus on and must fully understand each individual representation. But, if the representations are made too complex, then learners may not be able to isolate those dimensions that are shared between the representations. To aid learners in the integration of the shared information, it is important that those representations that present the overlapping information should be co-present.

5.1.3 Designing MERs with different processes

The final aspect of employing MERs to support different ideas and processes is when a designer aims to exploit the different computational properties of the alternative representations. For example, in some situations it may be appropriate to display tabular representations to emphasise order and patterns in numbers; in another, graphs may help to show the continuous nature of a phenomenon being examined.

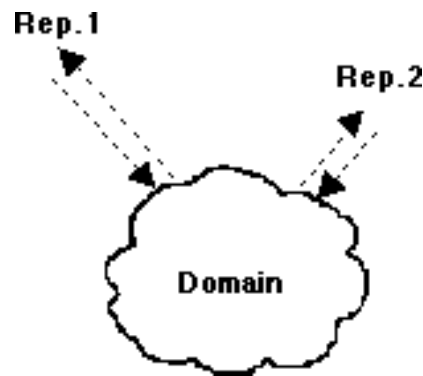


Figure 17. Using MERs to support different processes

The main focus for designing this use of MERs is the nature of the processes that each of the representations supports. In order for learners to benefit from these processes, again it is claimed that they should not be encouraged to co-ordinate the representations in the system. This claim is based on the recognition of the difficulties that learners have in translating between representations which appear dissimilar that was described in section 4.2.

Table 3. Design decisions of using MERs to support different processes

	Using MERs to support different processes
Redundancy	Full redundancy between representations
Similarity between Reps	Similarity determined by computational properties
Automatic Translation	Full automatic translation
Number of Representations	As many representations as task, learner or strategies require
Ordering of Representations	Limit co-presence. Display as task/learner/strategy demands

The majority of the design decisions follow from minimising the necessity for learners to translate between representations. Consequently, the representational system should be fully redundant so that a learner is not required to integrate information from different

sources. The computer rather than the learner should perform any translation that occurs and co-presence of the representations should be minimised to discourage unnecessary translation.

The advantages of combining representations with different computational properties can be found at the task, learner and strategies level (section 2.1.3). The number of representations and the order in which they are presented will depend upon these factors. For example, to successfully operate the device used in Bibby and Payne's experiments, three different tasks are performed and these tasks are each best supported by a different representation. It is not desirable or even possible to make a general recommendation that would account for all of these factors. Domain and learner specific knowledge is needed for these design issues.

5.2 Designing for constraining interpretation

The second broad class of functions that MERs serve is to constrain interpretations of a situation. One way that this may be achieved is to use a second representation to support interpretation of a more complicated, abstract or less familiar representation. The second type of constraint is when inherent properties of a representation can be used to support interpretation of another representation. Again, there are strong similarities between the proposed designs for both subclasses of constraining interpretation. In this case, both representations must be co-ordinated and this task should be made as easy as possible for the learner.

5.2.1 Designing MERs so that a familiar or concrete representation constrains interpretation of a second unfamiliar or abstract representation

Microworlds such as DM³ (Hennessy *et al.*, 1995) provide a simulation of a skater alongside a velocity-time graph. In such a situation, a common misunderstanding is that a straight line means no motion. This interpretation should be reconsidered by a learner when the simulation shows the skater still moving. In cases such as this, the second more familiar or concrete representation is not intended to provide new information about the domain, but to bridge understanding of the more complicated and unfamiliar representation.




Figure 18. Using MERs to constrain interpretation of a less familiar representation*

Table 4. Design aspects of using MERs so that a familiar representation constrains interpretation of a second unfamiliar representation

	Using MERs so a familiar/concrete rep. constrains interpretation of a second unfamiliar/abstract rep.
Redundancy	Full or subset redundancy between representations
Similarity of Representations	Maximise similarity
Automatic Translation	Full automatic translation
Number of Representations	One additional constraining representation
Order of Representations	Co-present

In order to achieve this use of MERs, the constraining representation should be as easy for a learner to understand as possible. Consequently, the first two learning demands should be kept to a minimum for this representation. This design can commonly be seen in the simulation environments that tend to include a concrete representation for this purpose. In addition, it is crucial that learners can easily co-ordinate the presented representations, otherwise the support for interpreting the unfamiliar representation will not occur. This suggests that representations that aid translation should be used. Hence, representations should be co-present and should be chosen to maximise the similarity between them (*e.g.* share the same labels, be in the same modality, have similar interfaces, *etc.*) Translation is also aided by full redundancy between the representations. Alternatively, if support for interpretation is required on only a limited number of dimensions of a complex representation, then the constraining representation could provide a subset of this information.

5.2.2 Designing MERs so that the inherent properties of the first representation constrains interpretation of a second representation

The second case of constraint between representations first introduced in section 3.3.2 is when constraints inherent in one representation affect the interpretation of another. This use

* The direction of arrow between the representations indicates the primary route by which translation between representations should occur. However, this is not intended to indicate that translation cannot occur in the opposite direction

of MERs is very similar to that of achieving constraint through the use of a familiar representation. However, in this case, it may not be possible to keep the learning demands of the constraining representations low (Figure 18). The example given in section 2.2.2, was the feedback provided to children on their solutions to coin problems by COPPERS. A property of the less familiar tabular representation (order irrelevance) constrains interpretation of the place value representations (which is order sensitive). Understanding order irrelevance is important if children are to recognise the commutativity in their solutions.




Figure 19. Using MERs to constrain interpretation of a second representation by exploiting the inherent properties of the first* .

Table 5 Design aspects of using MERs so that the inherent properties of a representation constrains interpretation of a second representation

	Using MERs so the inherent properties of one representation constrains interpretation of a second.
Redundancy	Full or subset redundancy between reps
Similarity of Representations	Similarity determined by computational properties, but aim to maximise wherever possible
Automatic Translation	Full automatic translation
Number of Representations	One additional constraining representation
Order of Representations	Co-present

If learners are to be able to benefit this intended use of MERs, then translation between the representations is crucial. However, the opportunity to maximise similarity between

* Figure 19 is very similar to Figure 18, the only significant difference is the amount of work needed to map between representations and between the rep 1 and the domain is likely to be greater.

representations may be reduced by the need to accommodate certain computational properties of the constraining representation. It is therefore important to make the task of co-ordinating the representations as easy as possible. To this end, the system should be designed to include automatic translation, co-present representations and either full redundancy or partial redundancy on the important dimensions.

5.3 Designing for Deeper Understanding

In section 2.3, three ways that MERs could lead to deeper understanding were introduced - through abstraction, extension and understanding the relation between representations. Although these are often seen as the most innovative aspects of multi-representational systems, there are fewer unambiguous empirical findings that can be drawn on to propose idealised designs for these functions of MERs. Furthermore, these uses of MERs are also seem most likely to co-occur with the other functions of multi-representational systems.

5.3.1 Designing MERs to promote abstraction

The defining criteria of abstraction that was introduced earlier (section 2.3.1) was as the process of re-organising knowledge at some **higher** level, through subtraction, reification or re-ontologisation (*e.g.* Giunchiglia & Walsh, 1991).




Figure 20. Using MERs to encourage abstraction

Table 6. Design aspects of using MERs to support abstraction

	Using MERs to support abstraction
Redundancy	Maximise redundancy between representations
Similarity of Representations	Unknown
Automatic Translation	Scaffolded translation
Number of Representations	Minimum number required to highlight invariances
Order of Representations	Co-present

Research reviewed earlier suggests that abstraction is a particularly difficult goal for learners to achieve (*e.g.* Schoenfeld; 1986, Sfard, 1991). Consequently, this use of MERs provides designers with hard choices. If users fail to translate across representations, then domain invariants are unlikely to be found. Experiments such as those reported with CENTS show that learners find translating over representations that are even superficially dissimilar to be difficult. However, in contrast to the cases of the constraining interpretation with MERs when representations also need to be reconciled, translation between representations should not be made too easy. If the alternative representations do not provide sufficiently different views on a domain, then the advantages associated with an abstraction are unlikely occur. For example, Dienes argues for perceptual variability in mathematics education - linking representations of a variety of formats. Balancing these two competing demands by identifying combinations of representations that can be co-ordinated by learners but which offer different perspectives is likely to be a far from easy task.

A similar worry concerns the role of automatic translation. It is known that learners are poor at recognising similarities and discrepancies between representation, (*e.g.* Borba, 1994). This is vital if domain invariants are to be uncovered. However, if the system performs all the translation activities for students, then they may not learn to co-ordinate the representations for themselves. In this case, it may be necessary to teach students to understand the relation between the representations first (section 5.3.3). In section 4.3, it was proposed that the best way to do achieve this understanding was to scaffold learners' understanding by dynamically reducing the support provided by the system as their competencies grows. Similarly, to maximise opportunities for learners to build cognitive links over representations, then representations should be co-present (*e.g.* Mayer, 1989). Ideally, fully redundant representations should be used as there is some evidence that suggests that this increases learners' abilities to reconcile representations that differ in format (Ainsworth *et al*, 1997a).

5.3.2 Designing MERs to encourage extension

In relation to learning with multiple representations, extension is considered to be the process of recognising that some aspect of a domain seen in a familiar representation can also be embodied by a new representation (figure 21). In this case, a learner is proposed to start with one known representation and is then taught to relate this representation to either an unfamiliar representation or a representation that has not previously been used for this purpose.




Figure 21. Using MERs to support extension

Table 7. Design aspects of using MERs to support extension

	Using MERs to support extension
Redundancy	Full redundancy between representations
Similarity between Reprs	Similarity determined by computational properties
Automatic Translation	Scaffolded translation
Number of Representations	One new representation at a time
Ordering of Representations	A new representation presented only when the existing representations are well understood

Teaching in this situation will be directed at helping learners to understand how the new representation relates to the familiar representation. It was claimed (section 4.3) that this is best achieved through scaffolding instruction of the relation between two representations. To minimise other learning demands, it is suggested that only one new representation is added at a time and that this should not occur until the existing representations are well understood. This can be determined by examining the extent to which a learner's behaviour has become co-ordinated over the representations already in use. To make the task of coordinating the representations easier, it is suggested that fully redundant representations are used, especially as the representations are likely to appear dissimilar to learners.

5.3.3 Designing MERs to teach the relation between representations

This use of MERs is close to that of extension and consequently the ideal design decisions look very similar. The difference is that rather than starting from one known representation and extending a learner’s knowledge from there, multiple representations are introduced pretty much simultaneously. This was illustrated in section 2.3.3 by referring to the CSCL system where three different representations of a skater’s movement are given at once in addition to the concrete simulation and with the QUADRATIC tutor that teaches 12-14 year-old pupils to relate an algebraic expression of the quadratic function to the area of a square.




Figure 22. Teaching the relations between MERs

Table 8. Design aspects of teaching the relation between representations

	Teaching the relation between MERs
Redundancy	Full redundancy between representations
Similarity of Representations	Similarity determined by computational properties
Automatic Translation	Scaffolded translation
Number of Representations	Two representations per relation to be learnt
Order of Representations	Co-present

As it is well known that learners find coming to understand the relation between representations difficult (*e.g.* sections 3.3, 3.4), then the design of the multi-representational system should help learners with this task by minimising the demands placed upon them. This can be achieved by presenting representations simultaneously, and maximising redundancy and, where possible, similarity between representations. Ideally, the number of new representations introduced at any one time should be limited. Again, it is suggested that an approach based on scaffolding the translation between representations will be the most successful way to help learners come to reconcile MERs .

5.4 Summary

Each of the eight functions of MERs proposed in section 2 of the paper can be seen to have a unique idealised design. However, the three classes of use for MERs that were proposed can be seen to have marked similarities in design. In particular, each class places a different emphasis on translation between representation. For different information and processes, translation between representations is considered unnecessary and consequently learners are not encouraged to learn how to relate representations in a system. When using MERs to constrain interpretation, it is crucial that the relation between representations is visible but this translation should be performed by the system. Finally, to achieve deeper understanding it is necessary for learners to achieve the cognitive linking of the representations. The best solution to this problem was suggested to be scaffolding a learner's understanding of the relation between representations.

These different functional designs are proposed as the basis for systematic investigation. It is hoped that they will be used to compare existing systems and experimental results and drive further exploration of what makes multi-representational learning successful. This may well lead to further extensions and clarification of this framework. One problem with this approach and with learning with MERs in general, is that often MERs are used for multiple purposes simultaneously. A system could teach the relation between representations in order to encourage abstraction or use representations with different computational properties and also develop extension. This suggests that a particular environment will often have to allow for multiple uses. Therefore, key features of a system may have to compromise on the fit between these proposed designs and the learning objectives or allow changes to the design in the life cycle of a system's use. To this end, authoring capabilities may be useful so that teachers and instructors can adjust the way that MERs are used within the system to the meet the changing needs of learners.

6.0 Conclusion

This paper has presented a framework for analysing the design of multi-representational learning environments. It consists of three elements, a functional taxonomy of MERs; a description of the learning demands associated with MERs and specification of design decisions unique to multi-representational learning environments. The final, more speculative, section combined these three elements in order to propose a set of idealised designs for each of the functions of MERs.

It was argued that multiple representations can serve many beneficial functions, especially when systems which employ them are designed to minimise learning demands. However, it should be pointed out that not all researchers are optimistic about the potential for multi-representational systems. In particular, Pimm (1995) warns that multiple linked representations may not be neutral. He suggests that one representation will come to

predominate and that by doing so it will no longer be viewed as a representation. Thus, meaning will not be associated with the relation between representations, but with the one dominant representation. Lowe (1997) also suggests that faced with multiple representations, learners often focus their attention on one dominant representation. He provides evidence that too often this focus is on the representation that is perceptually compelling rather than conceptually introducing. Finally, a number of studies by Sweller and colleagues (*e.g.* Chandler & Sweller, 1992; Kalyuga *et al*, 1998) have demonstrated that when information is presented in a number of representations rather than in a single representation, ‘split attention’ effects lead to increased cognitive load and less effective learning.

These arguments together with evidence presented in the rest of paper indicate that for learning with MERs to be successful, designers of software should carefully consider how they use multiple representations. The analytic framework developed in this paper is proposed as a further step towards the design, implementation and evaluation of effective multi-representational learning environments.

Acknowledgements

This work was supported by the Economic and Social Research Council at the ESRC Centre for Research in Development, Instruction and Training. A number of people have helped shape the ideas presented in this paper. My thanks to the members of the ESF-LHM taskforce on multiple representations, Pete Bibby, Peter Cheng, Rob Gaizauskas, Martin Oliver, Mike Scaife, Jean Underwood, Helen Woodiwiss and David Wood.

References

- Adelson, B. (1981). Problem solving and the development of abstract categories in programming languages. Memory and Cognition, 9, 422-433.
- Ainsworth, S. E. (1997). Designing and Evaluating Multi-Representational Learning Environments for Primary Mathematics PhD Thesis. (Also technical report number 47.). ESRC Centre for Research in Development, Instruction and Training, University of Nottingham.
- Ainsworth, S. E., Wood, D. J., & Bibby, P. A. (1996). Co-ordinating Multiple Representations in Computer Based Learning Environments. In Proceedings of the European Conference of Artificial Intelligence in Education, (pp. 336-342). Lisbon:
- Ainsworth, S. E., Wood, D. J. & Bibby, P.A. (1997a). Evaluating principles for Multi-Representational Learning Environments. In S. Vosniadou, K. Matsagouras, K. Maridaki-Kassotaki, & S. Kotsanis (Ed.), 7th European Conference for Research on Learning and Instruction, (pp. 500-501). Athens: Gutenberg University Publications.

- Ainsworth, S. E., Bibby, P. A., & Wood, D. J. (1997b). Information Technology and Multiple Representation: New Opportunities - New Problems. Journal of Information Technology for Teacher Education, 6(2).
- Ainsworth, S. E., Wood, D. J., & O'Malley, C. (1998). 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.
- Anderson, J. R., Conrad, F. G., & Corbett, A. T. (1989). Skill acquisition and the LISP tutor. Cognitive Science, 13, 467-505.
- Barwise, J., & Etchemendy, J. (1992). The language of first order logic: including the program Tarski's world 4.0. Center for the Study of Language and Information: CSLI lectures notes no. 34.
- Behr, M. J., Harel, G., Post, T., & Lesh, R. (1992). Rational number, ratio and proportion. In D. A. Grouws (Eds.), Handbook of teaching and learning mathematics (pp. 296-333). New York: Macmillan Publishing Company.
- Bibby, P. A., & Payne, S. J. (1993). Internalization and the use specificity of device knowledge. Human-Computer Interaction, 8(1), 25-56.
- Bibby, P.A., & Payne, S. (1996). Instruction and practice in learning to use a device. Cognitive Science, 20(4), 539-578.
- Blackwell, A. F., & Engelhardt, Y. (1998). A taxonomy of diagram taxonomies. Paper presented at the Third Conference of Thinking with Diagrams. Aberystwyth, August 1998.
- Borba, M. C. (1994). A model for student's understanding in a multi-representational software environment. In Proceedings of the 18th International Conference for the Psychology of Mathematics Education, vol. 2 (pp. 104-111). Lisbon.
- Boulton-Lewis, G. M., & Halford, G. S. (1990). An analysis of the value and limitation of mathematical representations used by teachers and young children. In Proceedings of the 14th International Conference for the Psychology of Mathematics Education, (pp. 199-206).
- Brna, P. (1996). Can't see the words for the tree: Interpretation and graphical representation. In I.E.E Colloquium on Thinking with Diagrams Digest no 96/010. London.
- Bruner, J. (1966). Towards a theory of instruction. New York: W.W. Norwood & Company.
- Case, R. (1985). Intellectual development. Orlando: FL: Academic Press.

- Chandler, P., & Sweller, J. (1992). The split-attention effect as a factor in the design of instruction. British Journal of Educational Psychology, *62*(2), 233-246.
- Cheng, P. C.-H. (1996a). Law encoding diagrams for instructional systems. Journal of Artificial Intelligence in Education, *7*(1), 33-74.
- Cheng, P. C.-H. (1996b). Learning qualitative relations in physics with law encoding diagrams. In G. W. Cottrell (Ed.), Proceeding of the Eighteenth Annual Conference of the Cognitive Science Society, (pp. 512-517). Hillsdale, NJ: LEA.
- Chi, M., T. H., Feltovich, P. J., & Glaser, R. (1981). Categorization and representation in physics problems by experts and novices. Cognitive Science, *5*, 121-152.
- Churchill, E. F., & Ainsworth, S. E. (1995). Making claims about teaching systems. In K. Nordy, P. Helmersen, D. J. Gilmore, & S. Arnesen (Ed.), Proceedings of Interact 95, (pp. 415-418). Lillehammer.
- Confrey, J. (1992). Function Probe©. Santa Barbara, CA: Intellimation library for the Macintosh.
- Confrey, J. (1994). Six approaches to transformations of functions using multi-representational software. In Proceedings of the 18th International Conference for the Psychology of Mathematics Education, vol. 2 (pp. 217-224). Lisbon.
- Confrey, J., & Smith, E. (1992). Revised accounts of the function concept using multi-representational software, contextual problems and students paths. In Proceedings of the 16th International Conference for the Psychology of Mathematics Education, vol. 1 (pp. 153-160). New Hampshire.
- Cox, R. (1996) Analytical reasoning with multiple external representations. Unpublished Ph.D. thesis, University of Edinburgh.
- 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, UK: NFER-Nelson.
- Dufour-Janvier, B., Bednarz, N., Belanger, M. & (1987). Pedagogical considerations concerning the problem of representation. In C. Janvier (Eds.), Problems of representation in the teaching and learning of mathematics (pp. 109-122). Hillsdale, NJ: LEA.

- Dugdale, S. (1982). Green globs: A micro-computer application for graphing of equations. Mathematics Teacher, 75, 208-214.
- Erhlich, K., & Johnson-Laird, P. N. (1982). Spatial descriptions and referential continuity. Journal of Verbal Learning and Verbal Behaviour, 21, 296-306.
- Fox, M. (1988) Theory and design for a visual calculator for arithmetic, CITE technical report no. 32, Institute of Educational Technology, Open University.
- Friedman, A (1993) Designing graphics to support mental models. In J.M. Spector, M.C. Polson, D.J. Muraida (Eds) Automating instructional design: Concepts and issues. Educational Technology Publications, New Jersey.
- Gentner, D. & Toupon, C. (1986) Systematicity and surface similarity in the development of analogy. Cognitive Science, 10, 277-300.
- Gick, M. L., & Holyoak, K.J. (1980) Analogical problem solving. Cognitive Psychology, 12, 306-355.
- Gilmore, D. J. (1996). The relevance of HCI guidelines for educational interfaces. Machine-Mediated Learning, 5(2), 119-133.
- Gilmore, D. J., & Green, T. R. G. (1984). Comprehension and recall of miniature programs. International Journal of Man-Machine Studies, 21, 31-48.
- Green, T. R. G. (1989). Cognitive dimensions of notations. In A. Sutcliffe & L. Macaulay (Ed.), People and Computers V. (pp 443-460). Cambridge: Cambridge University Press.
- Giunchiglia, F., & Walsh, T. (1992). A theory of abstraction. Artificial Intelligence, 57(323-389).
- Green, T. R. G., Petre, M., & Bellamy, R. K. E. (1991). Comprehensibility of visual and textual programs: a test of superlativism against the 'match-mismatch' conjecture. In Proceedings of Empirical Studies of Programmers: Fourth Workshop (pp. 121-146).
- Green, T. R. G., & Petre, M. (1996). Usability analysis of visual programming environments: a 'cognitive dimensions' framework. Journal of Visual Languages and Visual Computing, 7, 134-174.
- Halford, G. S. (1993). Children's understanding: The development of mental models. Hillsdale, NJ: LEA.
- 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.

- Janvier, C. (1987). Translation processes in mathematics education. In C. Janvier (Ed.), Problems of representation in the teaching and learning of mathematics (pp. 27-32). Hillsdale, NJ: LEA.
- Jones, S. (1998). Diagram representation: a comparison of animated and static formats (Cognitive Science Research Paper No. 493). School of Computing and Cognitive Sciences, University of Sussex.
- Kalyguga, S., Chandler, P., & Sweller, J. (1998). Levels of expertise and instructional design. Human Factors, *40*(1), 1-17.
- Kaput, J. J. (1987). Towards a theory of symbol use in mathematics. In C. Janvier (Ed.), Problems of representation in the teaching and learning of mathematics (pp. 159-196) Hillsdale, NJ: LEA.
- Kaput, J. J. (1989). Linking representations in the symbol systems of algebra. In S. Wagner & C. Kieran (Eds.), Research issues in the learning and teaching of algebra (pp. 167-194). Hillsdale, NJ: LEA.
- Kaput, J. J. (1992). Technology and mathematics education. In D. A. Grouws (Ed.), Handbook of teaching and learning mathematics (pp. 515-556) New York: Macmillan Publishing Company.
- Kaput, J. (1994). Democratizing access to calculus: New routes using old roots. In A. Schoenfeld (Eds.), Mathematical thinking and problem solving (pp. 77-156). Hillsdale, NJ: LEA.
- Laborde, C. (1996). Towards a new role of diagrams in dynamic geometry? In Proceedings of the European Conference of Artificial Intelligence in Education, (pp. 350-356). Lisbon.
- Larkin, J. H., & Simon, H. A. (1987). Why a diagram is (sometimes) worth ten thousand words. Cognitive Science, *11*, 65-99.
- Lesh, R., Post, T., & Behr, M. (1987). Representations and translations among representations in mathematics learning and problem solving. In C. Janvier (Ed.), Problems of representation in the teaching and learning of mathematics (pp 33-40). Hillsdale, NJ: LEA.
- Lohse, G. L., Biolsi, K., Walker, N., & Rueler, H. (1994). A classification of visual representations. Communications of the A.C.M., *37*(12), 36-49.
- Lowe, R. K. (1997). Effects of interactive animation in complex visual learning. In S. Vosniadou, K. Matsagouras, K. Maridaki-Kassotaki, & S. Kotsanis (Ed.), 7th European Conference for Research on Learning and Instruction, (pp. 181-182). Athens: Gutenberg University Publications.

- Mayer, R. E. (1989). Systematic thinking fostered by illustrations in scientific text. Journal of Educational Psychology, 81, 240-246.
- Meyer, J., Shinar, D., & Leiser, D. (1997). Multiple factors that determine performance with tables and graphs. Human Factors, 39(2), 268-286.
- Millwood, R. (1996). Issues in multimedia educational software design. Paper presented at the British HCI Group Usability and Educational Software Design London.
- Mulholland, P., & Domaigne, J. (1997). Teaching programming at a distance: The internet software visualisation laboratory. Journal of Interactive Media (<http://www-jime.open.ac.uk/97/1/>)
- Najjar, L. (1998). Principles of educational multi-media user interface design. Human Factors, 40(2), 311-323.
- Oberlander, J., Cox, R., Monaghan, P., Stenning, K., & Tobin, R. (1996). Individual differences in proof structures following multimodal logic teaching. In G. W. Cottrell (Ed.), Proceeding of the Eighteenth Annual Conference of the Cognitive Science Society, (pp 201-206). Hillsdale, NJ: LEA.
- Oliver, M. (1998) Visualisation and Manipulation Tools for Modal Logic. Unpublished PhD 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 Proceedings of the European Conference of Artificial Intelligence in Education, (pp. 357-363). Lisbon.
- Palmer, S. E. (1978). Fundamental aspects of cognitive representation. In E. Rosch & B. B. Lloyd (Eds.), Cognition and categorization (pp 259-303). Hillsdale, NJ: LEA.
- Pascual-Leone, J. A. (1970). A mathematical model for the transition rule in Piaget's developmental stages. Acta Psychologica, 32, 301-345.
- Petre (1993). Using graphical representations requires skill, and graphical readership is an acquired skill. In R. Cox, M. Petre, J. Lee, & P. Brna (Eds.), Proceedings of AI-ED93 workshop Graphical Representations, Reasoning and Communication, (pp. 55-58). Edinburgh.
- Petre, M., & Green, T. R. G. (1993). Learning to read graphics: Some evidence that 'seeing' an information display is an acquired skill. Journal of Visual Languages and Computing, 4, 55-70.
- Plötzner, R. (1995). The construction and co-ordination of complementary problem representations in physics. Journal of Artificial Intelligence in Education, 6(2/3), 203-238.

- Plötzner, R., Spada, H., Stumpf, M., & Opwis, K. (1990). Learning qualitative and quantitative reasoning in a microworld for elastic impacts. European Journal of Psychology of Education, *4*, 501-516.
- Pimm, D. (1995). Symbols and meanings in school mathematics. London: Routledge
- Pirolli, P., & Anderson, J. R. (1984). Learning to program recursion. Sixth Annual Cognitive Science Meetings, (pp. 277-280). Boulder, CO:
- Preece, J. (1983) Graphs are not straightforward. In T. R. G. Green, S. J. Payne, & G. C. van der Veer (Eds.), The psychology of computer use (pp 41-56) Academic Press: London.
- Resnick, L., & Omanson, S. (1987). Learning to understand arithmetic. In R. Glaser (Ed.), 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 (Eds.), Conceptual and procedural knowledge: The case of mathematics (pp. 225-265). Hillsdale, NJ: Lawrence Erlbaum Associates.
- 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 (Ed.), Advances in instructional psychology volume 4 (pp 55-175) Hillsdale, NJ: LEA.
- Schwartz, B., & Dreyfus, T. (1993). Measuring integration of information in multi-representational software. Interactive Learning Environments, *3*(3), 177-198.
- 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.
- Sfard, A. & Linchevski, L. (1994) The gains and pitfalls of reification – the case of algebra. Educational Studies in Mathematics, *26*, 191-228.
- Snow, R. E., & Yalow, E. (1982). Education and intelligence. In R. J. Sternberg (Eds.), A handbook of human intelligence (pp 493-586). Cambridge: Cambridge University Press.
- Spada, H., & Plötzner, R. (1997). Multiple mental representations: The case of qualitative and quantitative knowledge in physics. In S. Vosniadou, K. Matsagouras, K. Maridaki-

- Kassotaki, & S. Kotsanis (Ed.), 7th European Conference for Research on Learning and Instruction, (pp. 503). Athens: Gutenberg University Publications.
- Stenning, K., & Oberlander, J. (1995). A cognitive theory of graphical and linguistic reasoning: logic and implementation. Cognitive Science, 19, 97-140.
- Svendsen (1991). The influence of interface style on problem solving. International Journal of Man-Machine Studies, 35, 379-397.
- Sweller, J. (1988). Cognitive load during problem solving: Effects on learning. Cognitive Science, 12, 257-285.
- 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 (Eds.), 16th Annual Conference of the Cognitive Science Society, (pp 836-841). Hillsdale, NJ: LEA
- 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 (Ed.), 16th Annual Conference of the Cognitive Science Society, (pp 842-847). Hillsdale, NJ: LEA
- 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.
- Trowbridge, D. (1989) Graphs and Tracks©, Dept of Physics, University of Washington
- Underwood, J., & Underwood, G. (1987). Data organisation and retrieval by children. British Journal of Educational Psychology, 57, 313-329.
- Vernon, P. A. (. (1987). Speed of information-processing and intelligence. Norwood, NJ: Ablex.
- Vygotsky, L.S. (1978) Mind in society: The development of higher psychological processes. (M. Cole, V. John-Steiner, S. Scribner, E. Souberman, Eds. and Trans.). Cambridge, MA: Harvard University Press.
- Watson, J. M., Campbell, K. J., & Collis, K. F. (1993). Multimodal functioning in understanding fractions. Journal of Mathematical Behaviour, 12, 45-62.
- White, B. (1993). ThinkerTools: Causal models, conceptual change, and science education. Cognition and Instruction, 10(1), 1-100.

- Winn, B. (1987). Charts, graphs and diagrams in educational materials. In D. M. Willows & H. A. Houghton (Eds.), The psychology of illustration (pp. 152-198). New York: Springer-Verlag.
- Wood, D.J., Bruner, J.S. & Ross, G. (1976) The role of tutoring in problem solving. Journal of Child Psychology and Psychiatry, 17, 2, 89-100.
- Wood, D.J. & Wood, H.A. (in press) Help-seeking, learning and contingent computer-based tutoring. Computers & Education.
- Yerushalmy, M. (1991). Student perceptions of aspects of algebraic function using multiple representation software. Journal of Computer Assisted Learning, 7, 42-57.
- Zhang, J., & Norman, D. A. (1994). Representations in distributed cognitive tasks. Cognitive Science, 18, 87-122.

Appendix One: Examples of representations which display information concerning the age, gender, nationality and marital status of a population

	Nationality	French				Italian			
	Status	Single		Married		Single		Married	
Age	Gender	M	F	M	F	M	F	M	F
20s		5	17	12	9	13	8	17	12
30s		6	8	7	11	16	5	8	6
40s		8	12	6	7	6	8	12	7

Figure A1. Four dimensions of information in a tabular representation: age, gender, marital status and nationality




Figure A2 Four dimensions of information in a graphical representation: age, gender, marital status and nationality

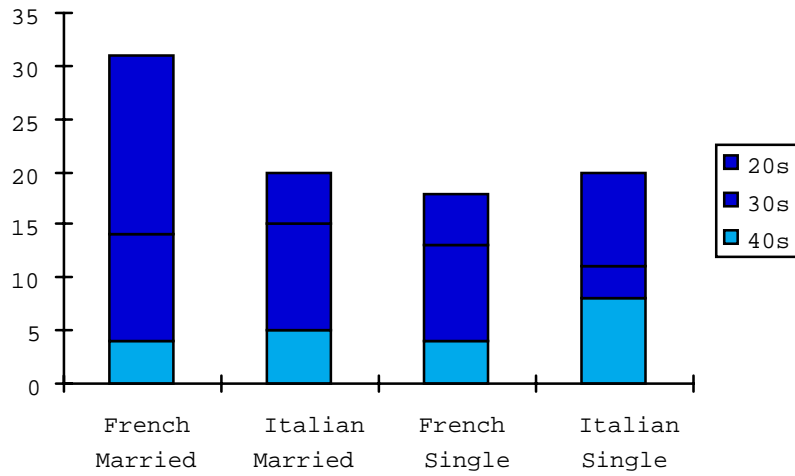


Figure A3 Three dimensions of information in a stacked histogram: age, marital status and nationality




Figure A4 Three dimensions of information in a bar chart: age, gender, marital status

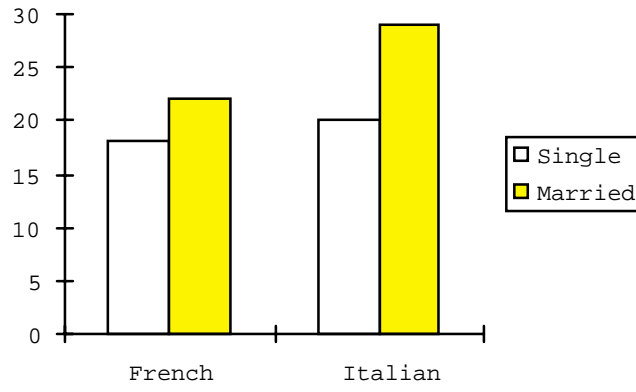


Figure A5 Two dimensions of information in a histogram: nationality and marital status

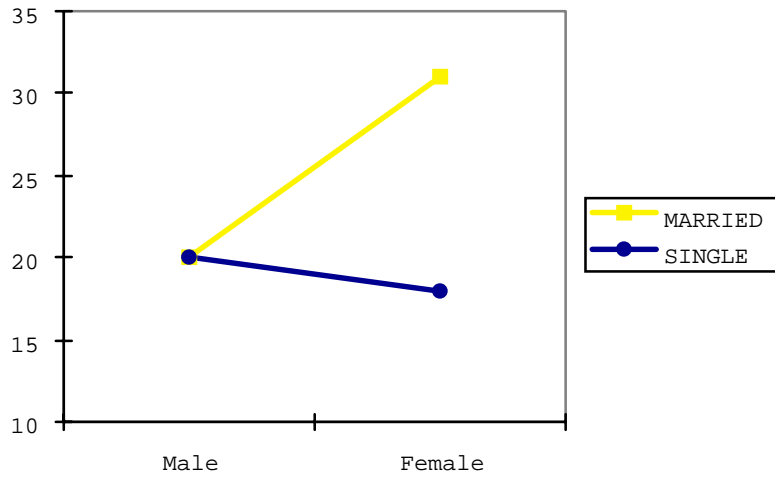


Figure A6 Two dimensions of information in a line graph: gender and marital status




Figure A7 One dimension of information in a line graph: nationality




Figure A8 One dimension of information in a histogram: marital status

Appendix Two: Examples of Multi-Representational Learning Environments

Name	No	Type	Redundancy	Order	Translation
MoLe	4	polygon world, modal world, natural language, predicate calculus	none - polygon & modal world, full - nat. lang. & pred. calculus.	learner control. Suggested route - modal world, polygon, nat. lang., pred. calculus,	none - polygon & modal world, signalled - pred. cal. & polygon world.
COPPERS (display)	4	pictorial, numerical, algebraic, mixed numerical & pictorial	full	pictorial, mixed numerical & pictorial, numerical, algebraic	automatic translation upon help
COPPERS (feedback)	2	table, numerical place value	full (per table row)	co-present	signalled
SkaterWorld (simulation)	9	1 from velocity, or distance or accel. time graph. All of pictorial simulation, tickertape, numerical display, net force indicator, force arrows, force controls	full - motion reps at a single point in time. full - force arrows & controls. full - net force and pairs of force arrows & controls	learner control - 1 from velocity, or distance or accel. time graph co-present - rest	automatic between force controls and other force reps. none - motion reps.
CENTS (feedback)	6	pictorial, numerical, <i>e.g.</i> splat wall, archery target, histogram, numbers	authorable between full, partial or non redundancy	2 co-present	none
TRM Microworld	3	table, graph, equation	full	learner control - 1 active rep at a time.	none
Point Grapher	3	table, graph, equations	full	co-present	automatic (from equation to table to graph)
Quadratic	2	iconic (picture of a square), equation	full	iconic to co-present	signalled
Blocks World	3	Dienes blocks, written maths, numbers	full	co-present	automatic
ReMIS-CL	8	numerical/equations, mass-velocity graph, 1d property diagram, velocity-velocity graph, polar graph, histogram, pictorial simulation, data trace.	full between all but data trace	choice of 1, 2 or 4 representations. Co-present with simulation	automatic between all but data trace.