In Lee, J. (1998) (ed.) Intelligence and Multimodality in Multimedia Interfaces: Research and Applications. AAAI. Press: Menlo Park, CA

How Can Interactive Multimedia Facilitate Learning?

Yvonne Rogers[1] and Mike Scaife
School of Cognitive and Computing Sciences
University of SussexBN1 9QH
Brighton

UK

Abstract

Educational and edutainment multimedia CD-ROMs, covering a diversity of topics, are on the increase. New titles are appearing every week: each offering ever more 'rich content' and 'genuine excitement' through 'exhilarating journeys' that are 'highly interactive' and which are 'lavishly illustrated, bursting with information'. In reality, however, many CD-ROMs have been poorly constructed, paying lip service to supporting the learning process. We argue that, if educational multimedia is to live up to its expectations and have a genuine pedagogical value, then much more focus is needed on how to support learning activities through designing effective interactivity. Our approach analyses how learners integrate information arising from different representations, through considering how external and internal representations are used in concert. From our theoretical analysis, we have constructed a set of design concepts that are instantiated as questions, issues and trade-offs for the design of interactive multimedia. To illustrate our approach, we present a series of demonstration spaces, varying in interactivity for teaching eco-concepts.

1. Introduction: more is more?

One of the main assumptions about multimedia is that the combination of graphics, video, sound, animation and text can provide better ways of presenting information than any of these media can alone. There is a general belief that 'more is more' and the 'sum is greater than the parts' (e.g. Lopuck, 1996). In addition, many claims have been made about the 'added value' gained from being able to *interact* with multimedia in ways not possible with single media (i.e. books, audio, video), such as easier learning, better understanding, more engagement and more pleasure. One of the main differences is that multimedia allows for rapid access to *multiple* representations of information. Many multimedia encyclopaedias[2] have been designed based on this *multiplicity*

---

[1] This paper was written whilst on sabbatical at the Computer Science Dept, Stanford University, USA.
[2] One of the best-selling educational CD-ROMs

principle, providing an assortment of audio and visual materials on any given topic. For example, if you want to find out about the heart, a typical multimedia-based encyclopaedia will provide you with:

• one or two video clips of a real live heart pumping and possibly a heart transplant operation
• audio recordings of the heart beating and maybe an eminent physician talking about the cause of heart disease
• various technicoloured schematic diagrams (static and animated) of the circulatory system sometimes with narration
• several columns of hypertext, describing the structure and function of the heart.

Judging by the numerous CD-ROM reviews appearing in the public forum (e.g. CD Source Line, CD-ROM on-line magazine, Anders CD-ROM survey) it is easy to see how pervasive the 'more is more' philosophy has become. Many view quantity and sophistication of multimedia information as important indicators of the value of an educational CD-ROM. For example, a typical review of Encarta'96, includes, phrases like, "it's packed with clever features, top-notch multimedia" and for The Way Things Work, "beautiful illustrations prompt wonderful animations...There is lots here...and it's all very accessible".

But what actually happens when users are given unlimited, easy access to multiple media? Do they systematically switch between the various media and 'read' all the multiple representations on a particular subject? Or, are they more selective in what they look at and listen to, adopting a more 'channel surfing' mode of interaction? Anyone who has explored a CD-ROM, knows just how tempting it is to play the video clips and animations, whilst skimming through accompanying text or static diagrams. The former are dynamic, easy and enjoyable to watch, whilst the latter are static, boring and difficult to read from the screen. For example, in an evaluation of a CD-ROM title on design (First Person) students consistently admitted to ignoring the text at the interface in search of clickable icons of the author, which when selected would present an animated video of him explaining some aspect about design (Rogers and Aldrich, 1995). Given the choice to explore multimedia material in numerous ways, paradoxically, users tend to be highly selective as to what they actually attend to. Hence, far from encouraging users to explore multiple representations, leading to the development of a more enriched understanding, multimedia environments may in fact induce more fragmented and superficial learning.

1.1 Limitations of current design of interactive multimedia

Part of the problem stems from the way many educational CD-ROMs have been designed. Typically, the different media are presented as separate entities at the computer interface. Text appears in one window whilst video clips and diagrams appear in other overlapping windows. The problem with this way of structuring information is that it is difficult, if not impossible, for someone to keep track of and integrate the separate sources. Simply, we do not necessarily have the cognitive resources to process information in this way. To do so we need to be able to switch our attention continuously within and between different parts of the various external representations in order to understand their content and relationship with each other.

Integrating multiple representations of information is probably much easier to achieve if they are physically adjacent to each other. This is something that paper-based educational materials can do relatively easily. Diagrams, depicting different abstractions, together with explanatory text can be presented simultaneously on the same page or sheet. Furthermore, external devices, like fingers, post-it notes and bookmarks, can be used by the reader to explicitly link information in the different representations. In contrast, multimedia representations are seriously lacking in their ability to support such interactions. Instead multimedia interactivity has been dominated by the 'tyranny of the button' (Hall, 1994). Many CD-ROMs have been designed to be used by simply pointing and clicking at menu options or hotspots on a display. The outcome is to bring up more information, be it more text, a new graphic or the running of an audio, video or animation clip. Further options are subsequently made available from which the user can choose to display other windows of information. Such 'point and click' style of interaction offers little more to the user than efficient page turning and channel-hopping capabilities.

It has to be said, though, that recently more effort has been made in designing multimedia applications to get users to do more 'activities' at the interface whilst pointing and clicking. For example, some CD-ROMs have been designed to provide electronic notebooks which the user can copy or type in their own material (e.g. What's the Secret?). Others include multiple choice quizs, geared towards testing the user's retention of facts, or puzzles, which the user has to solve by selecting and positioning different pieces in the right combination (e.g. Earth Explora, Kid's Encyclopaedia). Various simulation-type games have also been developed, where the user has to follow a set of procedures to achieve some goal for a given scenario. An example is the Cardiac Tutor built by Eliot and Woolf (1994) to teach students about cardiac resuscitation, where the students have to save patients by selecting the correct set of procedures in the correct order from various options displayed on the computer screen. Another innovation is the web-based Whole Frog Project, which allows students to explore the anatomy of a frog by selecting different options which result in different dissections occurring (Johnston, Nip and Logan, 1995). It also aims to teach students about the underlying imaging technology used to create the graphical structures. Another recent development is 3-D virtual environments, where pointing and clicking, effectively, are replaced by 'pushing' and 'pulling' actions resulting in the user moving through the environment. For example, CD-ROMs are now being designed to allow the user to travel through 3-D representations of an organism such as the human body (e.g. BodyWorks).

Whilst engaging the learner more with the information presented at the interface, these kinds of interactivities are still limited in how they support the learning process. For example, in an interview of an 11 year old who regularly played the popular simulation game, SimLife, Turkle (1995) reports how he played blindly, populating his world with creatures he likes rather than having any sense of purpose. When asked by Turkle why one species of animal has become extinct in his world he replied, "I don't know, it's just something that happens". When subsequently asked if he knew how to find out why this happened he replied no. When pressed further as to whether that concerned him, he answered, "No. I don't let things like that bother me. It's not what's important."

If educational multimedia is to live up to its expectations and have a genuine pedagogical value, the focus needs to be much more on how to support learning activities through designing effective interactivity.

2. Combining mutiple media: mix and match?

One approach to designing interactivity is to consider how best to combine multiple media in relation to different kinds of learning tasks e.g. when to use audio with graphics, sound with animations and so on. For example, Alty (1991) suggests that audio information is good for stimulating imagination; movies for action information, text for conveying details whilst diagrams are good at conveying ideas. From such generalisations it may then be possible to devise a presentation strategy for learning. This could be along the lines of: first, stimulate the imagination through playing an audio clip, then present an idea in diagrammatic form, then display further details about the concept through hypertext. More detailed guidelines have also been suggested along these lines. For example, Bernsen, (1993) has devised a set of guidelines, for helping in the selection of graphics in relation to the information that needs to be conveyed, e.g. "if what has to be displayed is a structural analysis of a complex abstract domain, then use network charts".

Such guidance about how best to select and combine multiple media, however, are often too general to help multimedia designers decide what *kinds* of audio, diagram, animation or text to use in relation to each other to represent the content of a particular topic in complementary ways. In particular, they offer little in the way of help for determining what to make explicit and salient about a domain in a particular representational form.

Another approach is to extrapolate the findings from the educational and psychological literature which has investigated how different representations of information might be integrated when learning. However, whilst there is a plethora of empirical studies that have investigated the advantages and disadvantages of learning from single or multiple representations of information, the findings are often conflicting. On the one hand, there have been many studies showing how presenting information in redundant modes can facilitate learning compared with using single media (for a review, see Levie and Lentz, 1982). Typically, these studies have compared the effects of simultaneous presentation of various combinations of media, such as spoken or written text with and without static diagrams or animations. For example, comprehension was found to be significantly better when information about how to operate a bicycle tire pump was depicted as an animation with concurrent narration, than when presented just as an animation (Mayer and Anderson, 1991).

On the other hand, there have been a number of studies that have shown that it can be more difficult for people to *integrate* written text with animations and diagrams. For example, when information about Newton's Laws of Motion were depicted as a series of animations with text explanations on a computer screen, Rieber (1989) found that subjects simply viewed the animations and then moved immediately onto the next screen of information without reading any of the accompanying text. In another study, retention of information was found to be significantly higher for paper-based presentations of biology texts and diagrams compared with computer versions (Reid and Beveridge, 1985). Here, the difference was attributed to the fact that the text and illustrations were displayed separately on the screen whilst being simultaneously available in the paper condition.

Whilst these studies can be revealing as to the effects of using different display formats and media on comprehension and retention of information, the findings are difficult to apply in a design context. Moreover, the criteria used in the design of the actual materials used in the studies is often largely absent. Hence, it is often unclear as to whether the various findings reported can be attributed to the form of the media, the mapping between the media and the underlying concept, the integration of the different media, or the kind of material being learned or the activity which the subjects have to perform.

3. An alternative approach for designing interactive educational multimedia

We argue that a more effective approach to the design of educational multimedia is one that focuses on how learners integrate information arising from different representations of the same and different information. This requires analysing how people learn to *read* and comprehend the significance of the content of different media and how this is assimilated with their current understanding of a domain. In turn, we need to ask: what is the best way of structuring different media, such that they convey the appropriate kind, level and abstraction of knowledge for a given domain?

Furthermore, we need to recognise the important role of constructing external representations (Reisberg, 1987), which is normally such an integral part of learning or problem-solving , e.g. highlighting text, making notes separately, in the margin or on overlaying post-its, re-representing text-based ideas in various canonical or idiosyncratic diagrammatic forms and sketching. We need to ask to what extent can they be supported, simulated or extended at the interface?

Taking these two concerns together, a central question is: how can we determine the most effective way of displaying and coordinating multiple representations at the interface whilst at the same time supporting the interactions and activities which the user should be able to control and do for themselves?

3.1 External cognition

In our research we have begun developing an alternative framework, which we call external cognition, for explaining how different representations are processed when performing different activities (Scaife and Rogers, 1996). Our starting point is to analyse the interactions between internal and external representations when used in concert (see also Larkin, 1989; Norman, 1993; Vera and Simon, 1993). Our rationale for adopting this stance, is that the process by which different external representations are used during learning or problem-solving is complex, involving the interplay between internal processes and different aspects of external representations at different stages of the cognitive task. For example, reading and abstracting knowledge from a diagram requires making connections between different elements of the display in a temporal sequence. In this respect, it doesn't make sense to model exclusively the putative internal processes that occur when using external representations to perform cognitive tasks; we need to explain how they are used together (cf. Glenberg and McDaniel, 1992).

One way of conceptualising the internal/external relationship is to identify cognitive properties that can be attributed to external representations which make them more or less easy to interact with (cf. Green's cognitive dimensions, 1989). 'Interact', in this context, refers to the various perceptual and cognitive processes that occur when external representations are used, adapted or constructed by the user in a given activity. These include searching, parsing, recognising, abstracting, re-representing, remembering and keeping track of different stages of a problem or activity. In our initial attempt to characterise these we have come up with the following: computational offloading, re-representation, graphical constraining and temporal and spatial constraining:

(i) computational offloading – This refers to the extent to which different external representations reduce the amount of cognitive effort required to solve informationally equivalent problems (Larkin and Simon, 1987).

(ii) re-representation – This refers to how different external representations, that have the same abstract structure, make problem-solving easier or more difficult (Zhang and Norman, 1994). It also refers to how different strategies and representations, varying in their efficiency for solving a problem, are selected and used by individuals.

(iii) graphical constraining – This refers to the way graphical elements in a graphical representation are able to constrain the kinds of inferences that can be made about the underlying represented concept (Stenning and Tobin, 1995; Stenning and Oberlander, 1995).

(iv) temporal and spatial constraining – This refers to the way different representations can make relevant aspects of processes and events more salient when distributed over time and space.

3.2 Cognitive characterisations and behavioural predictions

Together, these cognitive characterisations provide us with the starting point for a conceptual framework from which to consider why and how different kinds and combinations of external representations 'work' better than others for different cognitive activities. We can also use them to generate predictions about how effective different external representations are for different tasks.

Computational offloading

For example, we predict that the more computational offloading onto an external representation, the less mental effort is required by learners to grasp and apply the concepts being learned. However, there is a trade-off: we hypothesise that if there is too much computational offloading, whereby the solution required can be easily 'read-off' from the external representation, then the learners will only superficially process the content of the information. Thus an optimal level of computational offloading is needed to promote effective learning.

To test this we can vary the amount of inferencing required by the learners to solve a problem in a particular domain by using different forms of graphical representations (static, animated) which have different kinds of interactivity. The latter are operationalised in terms of level of user control, extent of annotating, amount of feedback and level of complexity of domain knowledge. The effects of varying these indices of interactivity can then be quantifiably and qualitatively measured in terms of retention of information, accuracy of solving problems and transfer of problem-solving strategies learned to similar problems.

3.3 Cognitive characterisations and design concepts

The cognitive characterisations also provide a framework from which to construct design concepts to frame questions, issues and trade-offs. By design concepts, we are refering to the issues which designers need to consider when developing interactive multimedia. The kinds of questions which the concepts raise, range from interaction concerns about how to achieve a particular goal to pedagogical concerns about whether certain designs are effective. For example, in the four design concepts outlined below (see Scaife and Rogers, 1996, for more details) the first two provide a statement of the concept followed by a question that asks how can this be achieved, whilst the second two provide a statement of the concept followed by a question of how this could improve learning.

• explicitness and visibility – how to make more salient certain aspects of a display such that they can be perceived and comprehended appropriately. *Example question: How can we use the visual and auditory properties of external representations to make explicit processes that are difficult to grasp because they are not perceptually obvious?*

• cognitive tracing – what are the best means to allow users to externally manipulate and make marks on different representations.

Example question: How is it possible to allow users to leave annotations and marks in dynamic representations like animations and 3-D virtual environments?

• ease of production – how easy it is for the user to create different kinds of external representations, e.g. diagrams and animations
Example question: Does an understanding of how to create external representations (diagrams, animations) enable people to have a better understanding of knowing when they are most effective to use?

•combinability and modifiability – how to enable the system and the users to combine hybrid representations, e.g. enabling animations and commentary to be constructed by the user which could be appended to static representations
Example question: Does the ability to create heterogeneous representations provide the user with a better understanding of the domain knowledge?

3.4 From design concepts to design parameters

At a finer level of description, the design concepts are realised in terms of particular parameters that can be implemented at the interface, such as the type of media, the kinds of navigation aids and use of colour. A key concern at this level is to consider how the different parameters can be coordinated effectively with each other. For example, when thinking about how to make certain aspects of different representations more explicit, at the relevant stages of learning, we could display a diagram with accompanying narration, that directs the user towards the salient parts of the diagram. Such *cueing* could help the novice also learn how to 'read' the diagram types for that domain through focusing their attention on the relevant aspects that need to be visualised (i.e. seen in relation to each other in order to make an inference) at a specified time. Redundant visual coding could also be used to further *constrain* the way the information should be interpreted. For example, flashing arrows could appear at relevant times, to highlight the elements of the diagram that are being referred to in the narration.

4. Domain knowledge

In addition to describing design issues in relation to our cognitive characterisations we also consider how we can use different representations to convey more effectively the domain knowledge that needs to be learned. One of our primary concerns is domain knowledge that has consistently proven to be difficult to learn, using existing single 'non-responsive'[3] media (e.g. textbooks, videos and audio). For example, it is well known that children (and adults) have difficulty understanding and using correctly the formal notations and abstract diagrams that are used in the domains of physics, biology, maths and logic. This is particularly problematic, since the formal representations are not merely explanation aids, but are an essential part of the domain knowledge itself (e.g. Anzai, 1991). Why is it that they are so difficult to understand and construct?

Part of the problem of learning how to use formal representations is that they are used to represent complex abstract concepts. Often these are invisible processes, such as electricity, acceleration, and various kinds of cycles and trends. Hence, they do not resemble or map onto the concept to which they refer in an intuitive way – as a picture or animation of an object in motion can do. For example, Schwartz (1995) points out how a Cartesian graph of the Gross National Product (GNP) does not look like it but covaries with it: the line used in the graph does not explain what GNP is but conveys how it changes across time. Similarly, schematic diagrams used in ecology to show feeding relationships between different organisms convey the direction of flow between the organisms rather than the process taking place. Here, children have a hard time trying to relate these formal abstractions (e.g. chains and webs) to what they already know happens in the physical world. On the other hand, they have no problem relating to the physical instantiations of these processes, such as a pigeon eating a slug or a slug eating a lettuce. Hence, whilst is easy to understand that certain organisms eat other organisms, it is much more difficult to understand them at a higher level of abstraction, for example in terms of primary, secondary and tertiary consumers.

Our approach to the abstraction problem has been to consider how we could design different kinds of interactive representations which could convey more explicitly the relations between different levels of abstraction of information. By

[3] By non-responsive, we mean that the media has a fixed presentation. For example, a book can be interacted with in different ways, e.g. read from cover to cover or flicked through, but the format is constant. The text or the diagrams cannot change into something else in the way in which hypertext or animated diagrams can be designed to do at the interface.

designing dynamic interlinked representations that co-vary in their abstraction and interactivity it may be possible to provide children with a more effective way of understanding and reconstructing the formal notations used to describe the concepts. We can also vary the level of computational offloading. For example, when introducing an abstract concept, we could depict a simple animation of a concrete instantiation of it. Here the computational offloading would be high; the user is required just to observe the concrete events that occur. Later on, when we want to teach the concept at a higher level of abstraction, we can decrease the computational offloading so that more cognitive processing is required by the user. For example, we could provide them with an empty diagram which they have to complete by combining the appropriate elements in the correct places.

In conjunction, we can exploit the capabilities of multimedia to explicitly and dynamically show how one kind of representation can be re-represented at both higher and lower levels of abstraction. For example, the effect of making changes to elements in an ecosystem conveyed in one kind of representation could be simultaneously depicted in another kind of representation. Two (or possibly more) external representations, varying in abstraction (e.g. a schematic diagram, an animation, a 3-D world), could be displayed on the same screen. The user could also be given the ability to make changes (correct or incorrect) to the elements in the eco-system in any of the representations, to which the corresponding elements in the other representations would change accordingly.

4.1 An example: PondWorld

In our research project, called eco-i[4], we have begun developing demonstration spaces that depict eco-concepts as a progressive series of interactive representations that have been designed to vary in their level of abstraction and interactivity. For our first implementation, we have begun developing a series of PondWorld demonstration spaces to convey the abstract concepts of food webs and food chains. The target group is, initially, for 11-14 year olds. The first demonstation space has been designed to show an animation of a simple pondworld ecosystem, with fish predators eating water beetles, water beetles eating tadpoles and tadpoles consuming weeds. Users can easily add or remove

[4] Eco-i stands for "Explaining External Cognition for Designing and Engineering Interactivity" and is a project funded by the Cognitive Engineering Initiative.

organisms so that they can focus their attention on what happens between different combinations of the organisms. Here the level of computational offloading is high: users have only to observe the animation. The kind of interactivity has also been designed to engage the learner in guiding their attention to different relationships. Narration is also available to describe what is going on in the animations.

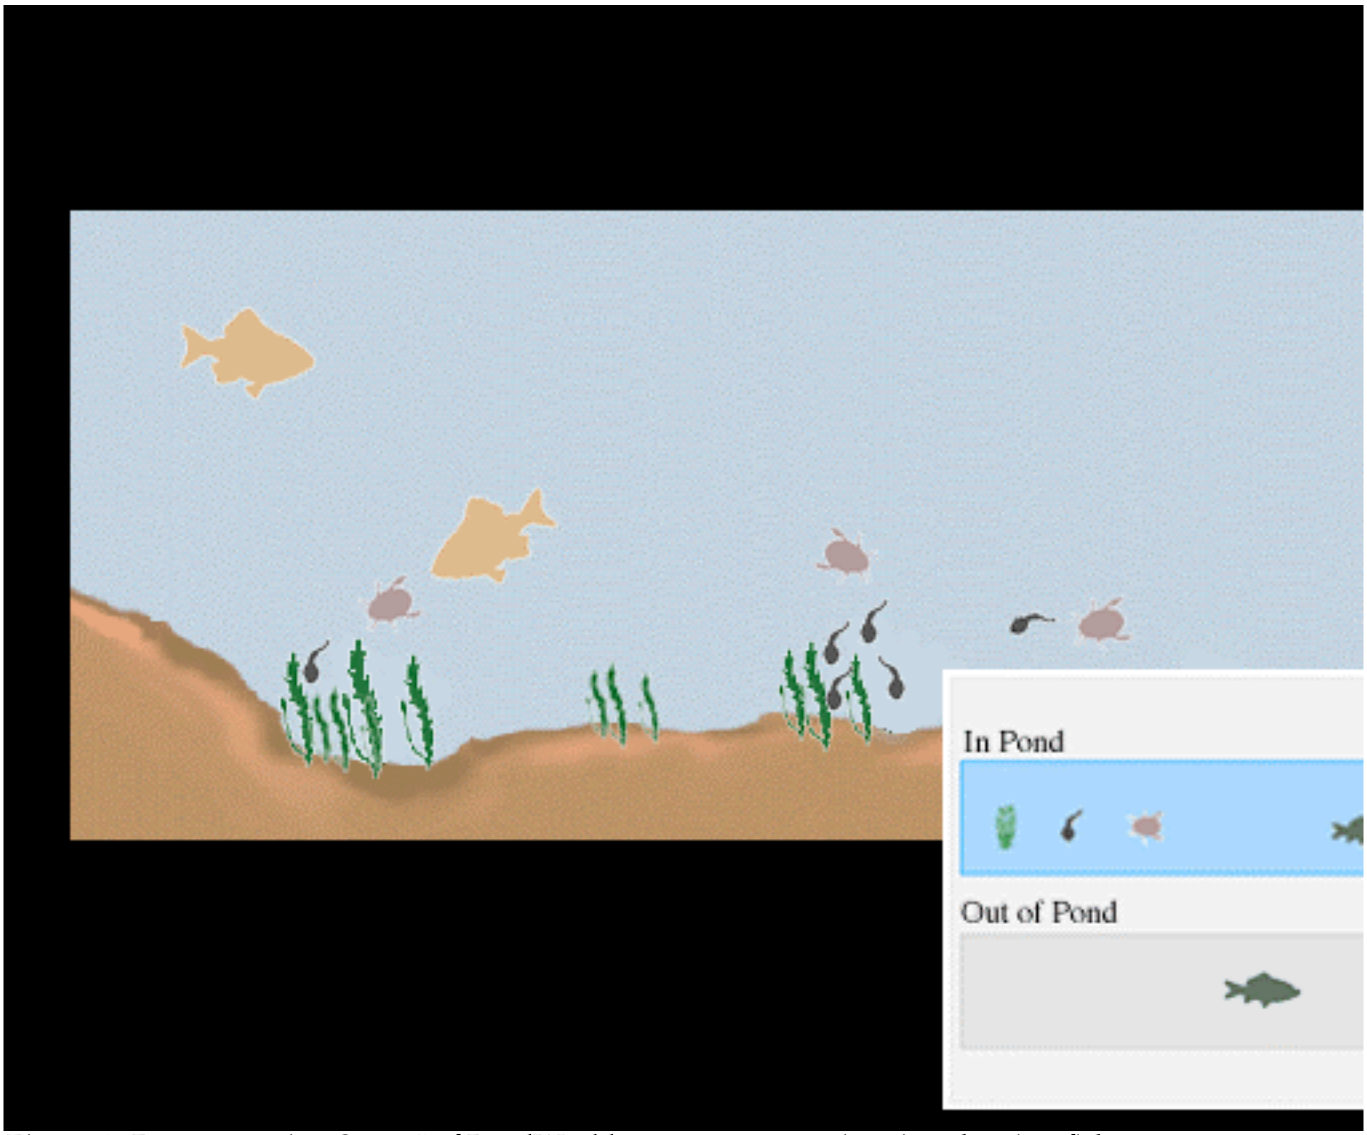

Figure 1. Demonstration Space I of PondWorld ecosystem: an animation showing fish predators eating water beetles, water beetles eating tadpoles and tadpoles consuming weeds. Organisms can be added or removed from the pond through turning them on or off using the accompanying palette.

The second demonstration space is designed to allow the learner to construct their own ecosystem which is analogous to the one they have just interacted with. They are provided with a jamjar and a palette of organisms which they can select and place in the jamjar. They also have the option of removing the organisms using a simulated fishing net. Having placed different combinations of organisms in to the jamjar the user can animate them and observe what happens. Running commentary can also be switched on at various points to focus the learner on any salient changes in the populations they have created (e.g. "watch the small fish. See how they gobble up the snails.") In this demonstration space, the learner is actively engaged in constructing their own mini ecosystem, whilst the system visualises for them what happens to the organisms when combined in such ways. At certain significant points, e.g. following the extinction of a population of tadpoles, the system will pop up a question, like "what will happen to the beetles now?" Here temporal constraining is important in focusing the learner's attention on salient aspects of the relationship between different populations. The computational offloading is kept relatively high, but with more problem-solving activities involved.

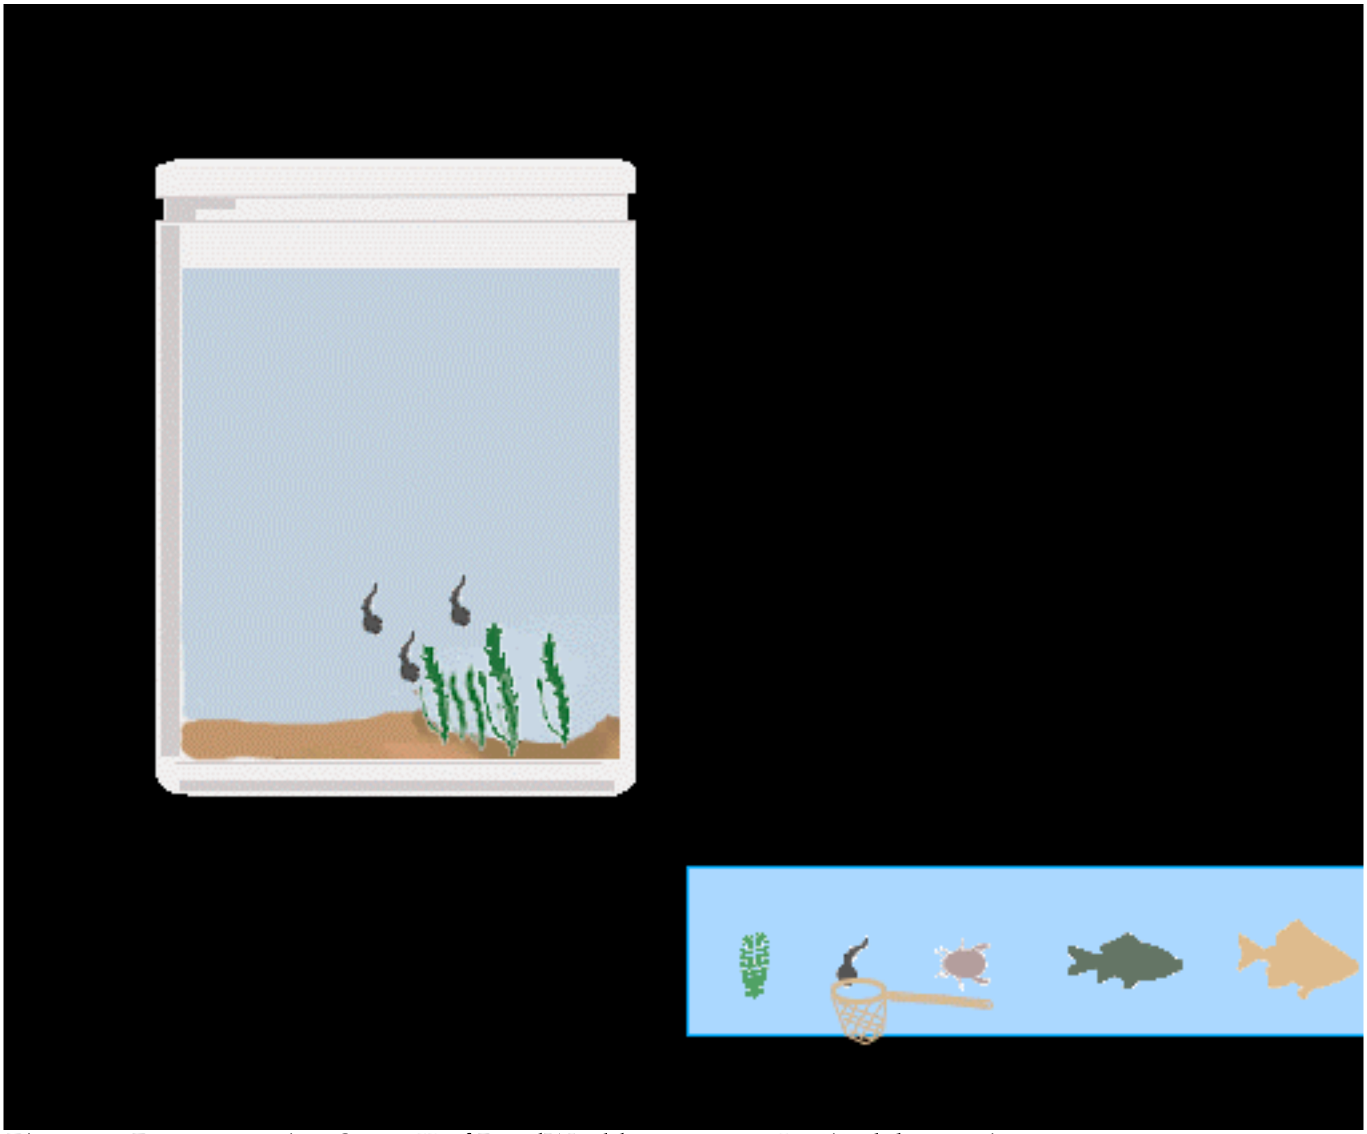

Figure 2. Demonstration Space II of PondWorld ecosystem: Jamjar lab experiment metaphor. The learner is required to construct their own jamjar ecosystem analogous to PondWorld by selecting organisms from the palette and placing them in the jamjar. The organisms also can be removed with the fishing net. The jamjar ecosystem then can be animated for different combinations of organisms.

In the next demonstration space the learner is presented with two adjacent displays: a canonical food web diagram and a concrete instantiation of it. The former is an abstraction of the latter and our goal is to make this explicit to the learner. We do this by integrating the two representations: the organisms in the animation are designed to behave in relation to the abstract feeding relationships depicted in the food web. The learner can 'turn on' different combinations of feeding relationships in the food web (e.g. the link between the weeds and tadpoles) and observe the weeds being eaten by the tadpoles. Here the computational offloading is high; the learner is required to select and observe the outcomes of their intervention.

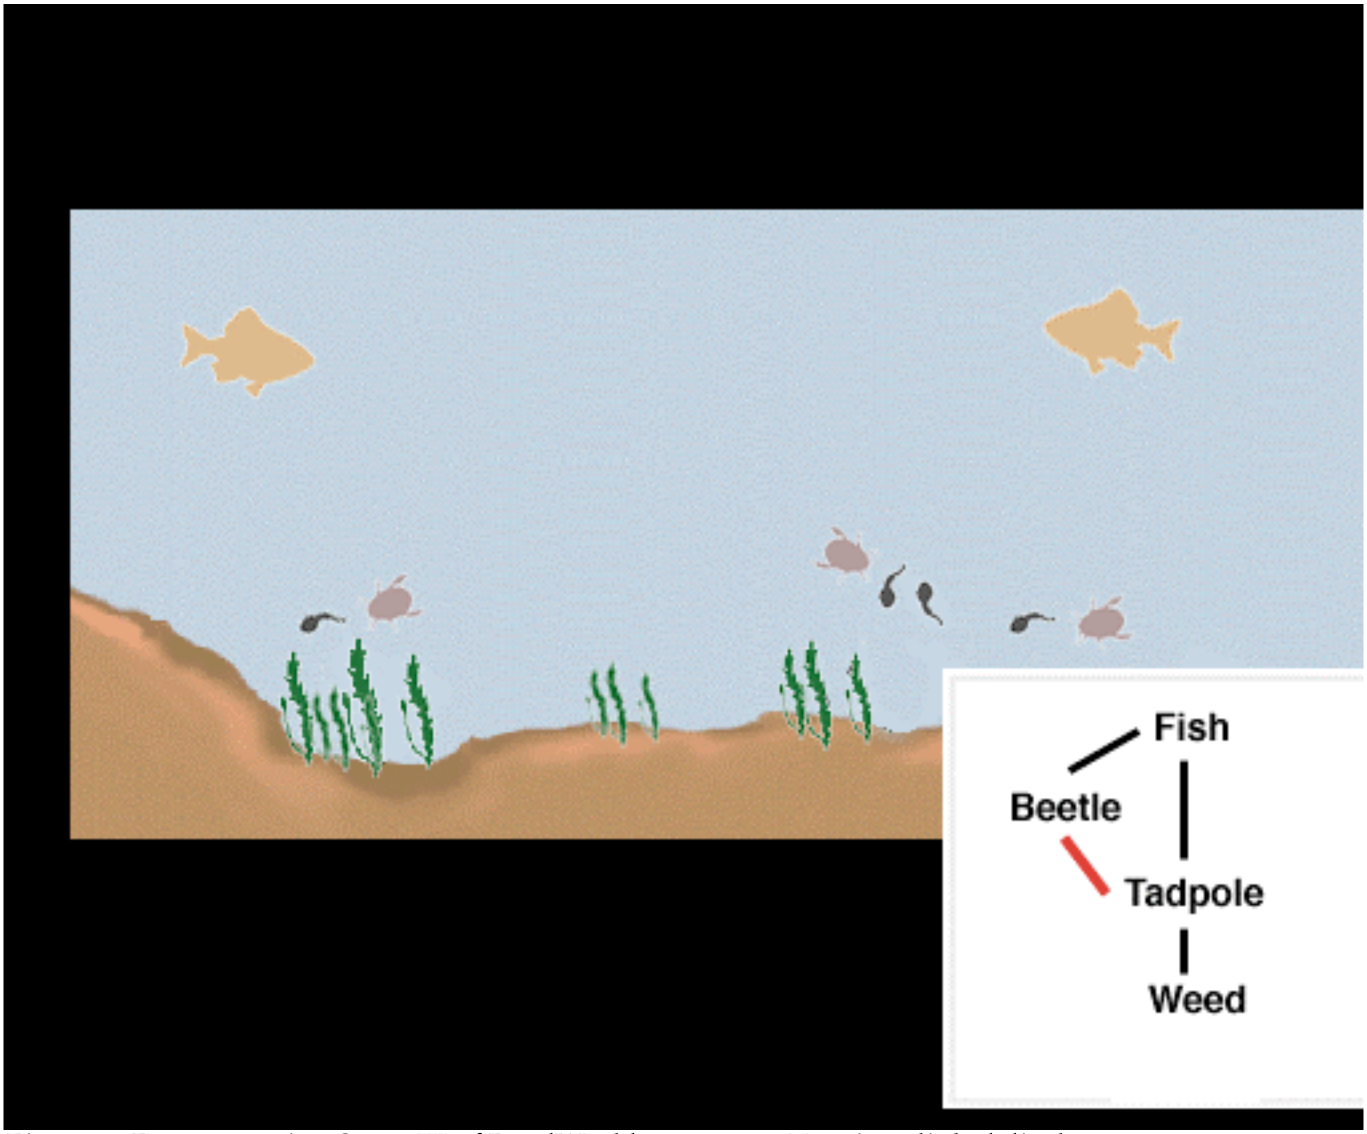

Figure 3. Demonstration Space III of PondWorld ecosystem: Two interlinked displays are shown: a canonical food web diagram and a concrete animation of it. The learner can 'turn on' different combinations of feeding relationships in the food web (shown by red line) and observe the weeds being eaten by the tadpoles in the adjacent animation

In the next demonstration space we show a similar split screen set-up with two linked displays. This time the food web is empty and the learner is required to create their own food webs from the palette of organisms presented in a previous demonstration space. Again, an important feature of this demonstration space is that the effects of creating a food web as an abstract representation can be directly viewed in the adjacent concrete animation. The environment is also set up so that the learner can create incorrect food webs and observe the outcome. For example, they can place the weeds above the tadpole in the food web and observe the weeds eating the tadpole. As children of that age know well that weeds do not eat tadpoles the animation can provide them with a fun way of showing them that they have placed the organisms the wrong way round in the abstract food web diagram.

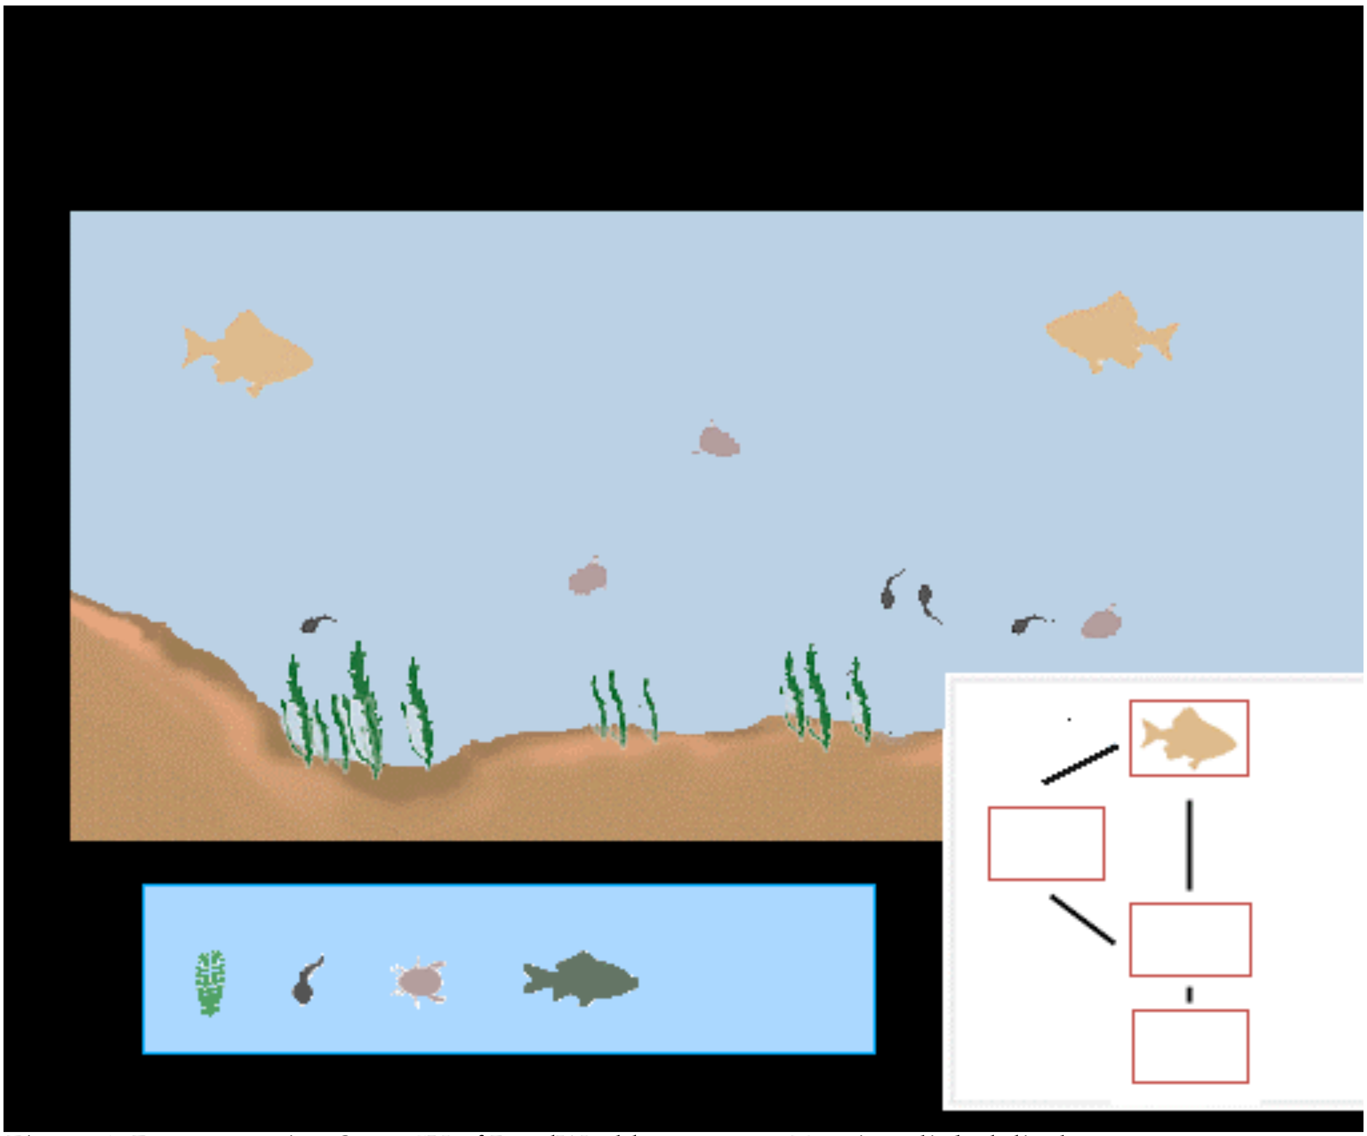

*Figure 4: Demonstration Space IV of PondWorld ecosystem: Two interlinked displays are shown: a canonical **empty** food web diagram and a concrete animation. The learner has to create their own abstract diagram by selecting organisms from the palette. These can also change to text labels in the diagram. Different food web diagrams can be seen in action in the adjacent animation.*

In subsequent demonstration spaces we intend to implement further eco-concepts, varying in their level of abstraction, using similar combinations of representations and interactivity. We also intend to provide other kinds of interactivity (e.g. annotating, creating animations) depending on the kind of learning activities we are supporting. As with previous demonstration spaces, for each new one we consider how to characterise it cognitively, i.e. how much cognitive offloading to provide; how to make salient certain aspects of the representations which are considered important at that stage of learning about the concept using temporal and graphical constraining; how best to re-represent a concept and so on. Each new demonstration space is built on the assumption that the learner has already integrated the knowledge presented in the previous demonstration spaces.

5. Discussion and conclusions

In some ways, our idea of providing various forms of interactivity that allow learners to explore the outcome of their actions is similar to one of the pedagogical goals behind many of the simulation microworlds built in the 80s and more recent modelling learning environments. For example, ARK (alternate reality kit) developed by Smith (1986) was designed to allow children to explore what happens to the behaviour of objects when physical laws are violated. A variety of objects were provided at the interface whose properties changed in response to user interaction. Objects like balls, could change their velocity and position in accordance with changes made by the user to 'gravity' and 'motion' sliders. Likewise, Soloway et al's (1996) more recent modelling environments have been designed to enable children to build models, such as stream ecosystems, using physical objects which they define relationships between. The system can then show these relationships as both textual and graphical representations, e.g. the Model-It software displays the relationship 'decreases by less and less' as a static graphical representation. (see figure 5).

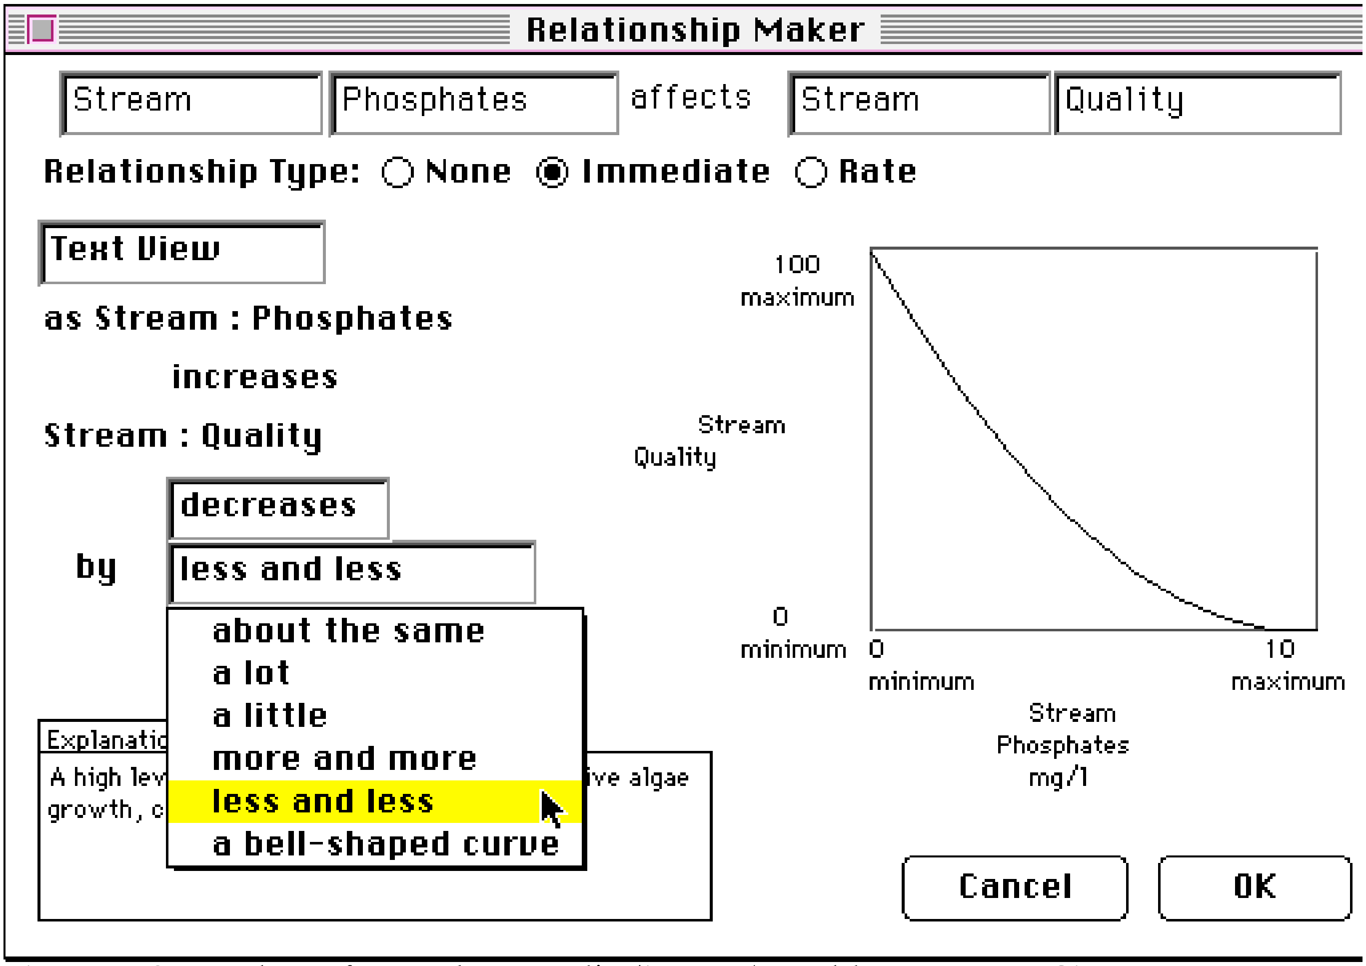

Figure 5. Screen dump from Jackson et al's (in press) Model It program. Given a qualitative definition, e.g. "as stream phosphate increases, stream quality decreases by less and less" the Model-It software translates it into a quantitative visual representation as shown by the graph.

Where our approach differs from either the simulation or modelling approaches, is on its emphasis on using appropriate kinds of interactivity and external representations at different stages of learning, per se. In particular, our focus is very much on making the abstractions that are part of a domain's knowledge more explicit. We should also stress that we are not proposing an alternative way of learning, but that we can use interactive multimedia more effectively to explain and convey abstract concepts, through varying the way the learner can visualise and interact with the different representations.

We base this on an evolving view of the nature of external representations and the ways that they function, cognitively, in relation to internal models. The central idea of 'bridging the gap' between abstraction and reality which the prototypes embody is, therefore, not just grounded in the *intuition* that learning must be easier in smaller steps. Rather it derives from a theoretical view about how to describe the properties of building representations.

Finally, we would like to mention Stenning and his colleague's work at Edinburgh (e.g. see this collection) where they, too, have designed and used a range of external representations for teaching the abstract domain notation in logic necessary to deduce proofs in predicate calculus. Their research has investigated how different kinds of graphical representations (e.g. Euler's circles and 3-D shapes like cubes), varying in their expressiveness, can facilitate learning. In one study, they found that performance on reasoning improved when students learned using Hyperproof[5] (Cox et al, this volume). One of the reasons for this is attributed to the fact that Hyperproof has been explicitly designed to facilitate the translation of information between graphical and sentential representations, through displaying two interlinked graphical and calculus 'worlds' at the interface (see Figure 6).

[5] Hyperproof was designed by Barwise and Etchemendy (1994) as an interactive multimedia environment for teaching first order logic.

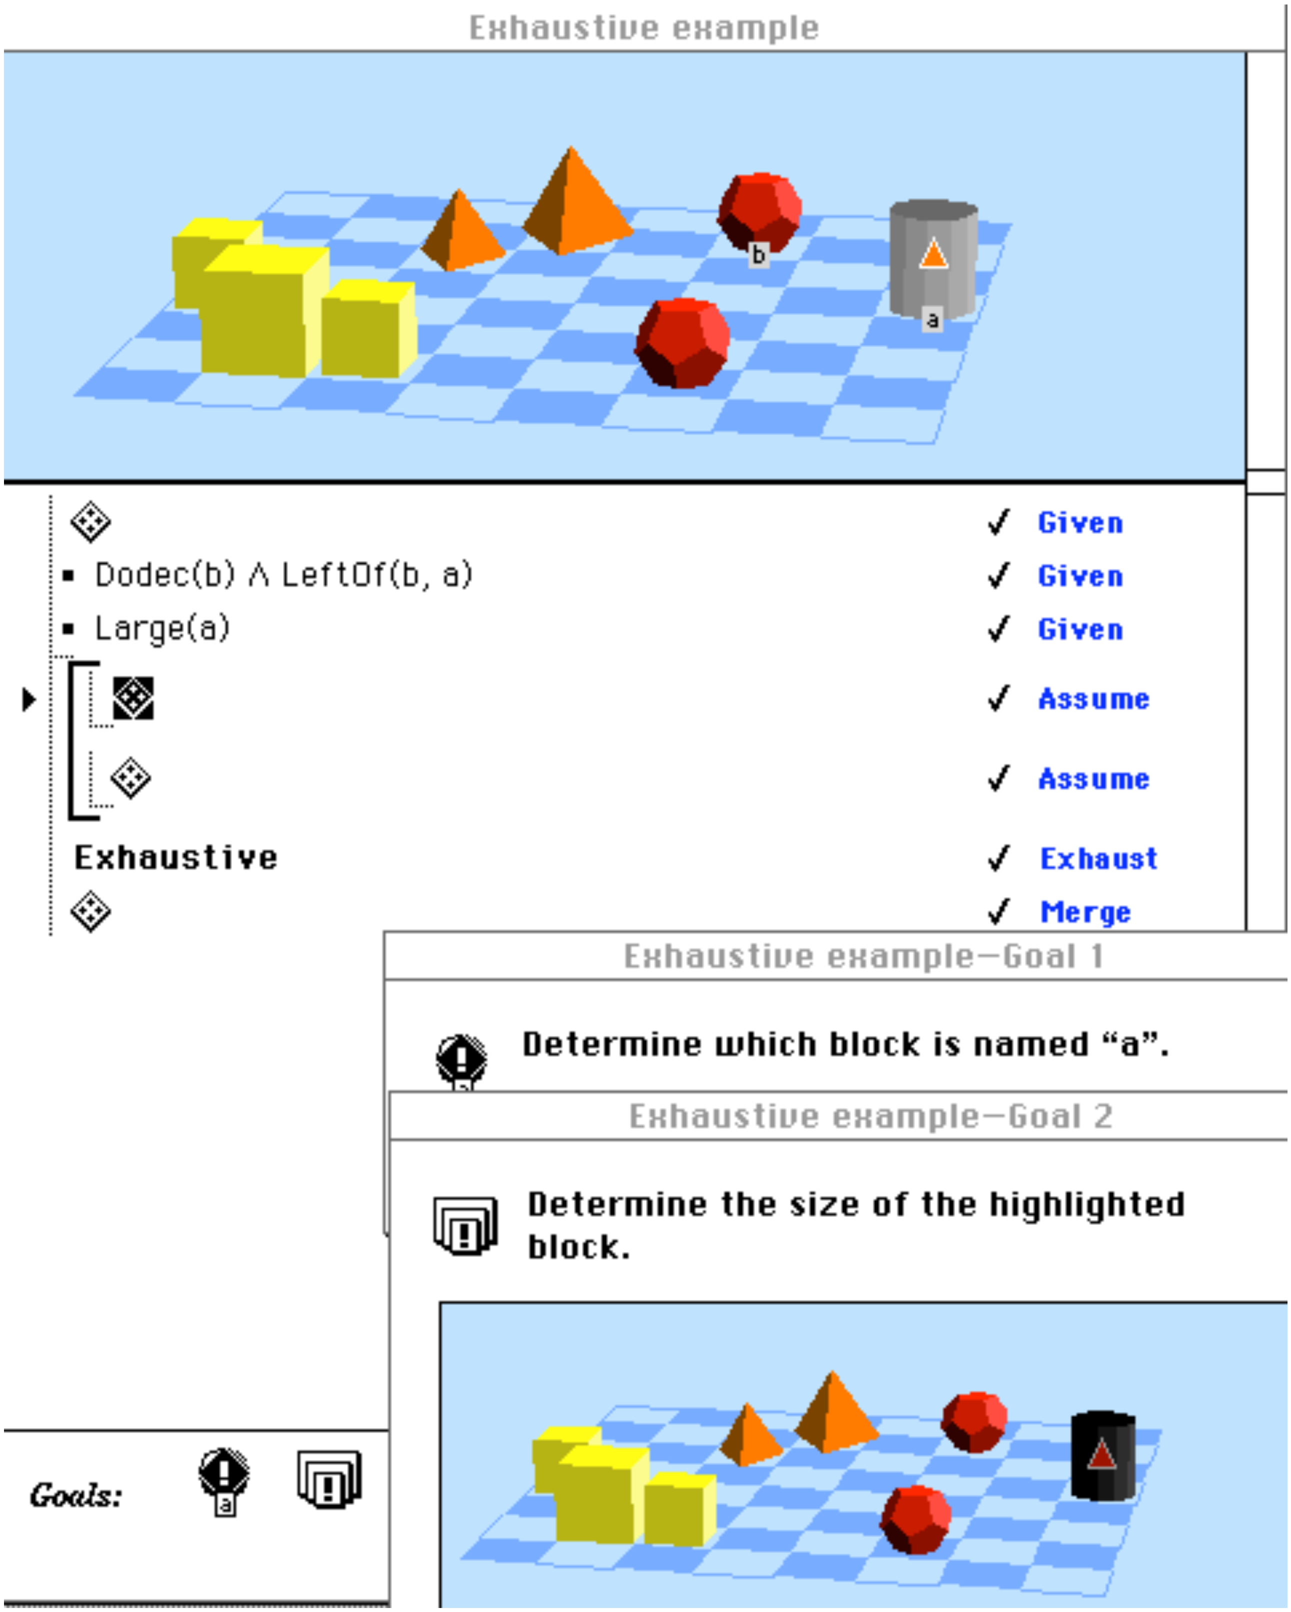

Figure 6. The graphical and calculus worlds of Hyperproof. From Cox et al.

This idea corroborates our analysis about how the translation between different representations, both internally and externally is central to effective learning. They conclude by citing Barwise (1993): "efficient reasoning is inescapably heterogeneous (or 'hybrid') in nature" and suggest, accordingly, that we should be exploiting interactive multimedia to allow users to generate their own heterogeneous representations. Our way of more effectively exploiting the potential of interactive multimedia is to design interactive demonstration spaces that can explicitly and dynamically co-vary the effects of changing aspects in one form of representation in another. More generally, our goal is to design interactive learning environments, which support the efficient interaction of internal and external representations during cognitive activities.

Acknowledgements

The authors gratefully acknowledge support from the ESRC Cognitive Engineering Programme, project L127 30 1001 66. Yvonne Rogers would also like to thank the People, Computers and Design (PCD) Group for their support whilst on sabbatical at Stanford University, Computer Science Dept. We also thank Matt Davies for his prototype designs for PondWorld.

References

Blue coding: Paper-based journal articles and books

Alty, J.L. (1991) Multimedia – What is It and How do we Exploit It? In D. Diaper and N.Hammond (eds) *People and Computers IV,* CUP: Cambridge.

Anzai, Y. (1991) Learning and use of representations for physics expertise. In K. A. Ericsson & J. Smith, Eds. *Towards a general theory of expertise.* pp. 64-92. Cambridge: Cambridge University Press.

Barwise, J. (1993) Heterogeneous reasoning. In G. Allwein & J. Barwise (eds.) *Working Papers on diagrams and logic,* Indiana University Logic Group Preprint No. IULG-93-24, May 1-13.

Bernsen, N.O. (1993) Matching Information and Interface Modalities: An example Study, *WPCS-92-1,* Roskilde University.

Cox, R., Stenning, K. and Oberlander, J. (this volume) The effect of interactive multimedia interfaces upon representation selection.

Eliot, C and Woolf, B. (1994) Reasoning about the user within a simulation-based real-time training system. In *Proceedings of 4th International Conference on User Modeling,* Mitre Corp., Bedford, Mass.

Glenberg, A.M. and McDaniel, M. A. (1992) Mental Models, pictures and text: Integration of spatial and verbal information. *Memory and Cognition,* 2(5), 458-460.

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

Hall, W. (1994) Ending the Tyranny of the Button. *IEEE Multimedia,* 1, 1, pp. 60-68.

Jackson, S. L., Stratford, S. J., Krajcik, J. S., Soloway, E. (to appear) Making Dynamics Modeling Accessible to Pre-College Science Students. *Interactive Learning Environments.*

Larkin, J. H. (1989) Display-based problem solving. In: D. Klahr & K. Kotovsky, Eds. *Complex Information Processing: The Impact of Herbert A. Simon.* Hillsdale, NJ: Lawrence Erlbaum Associates. pp. 319-341.

Larkin, J. H., & Simon, H. A. (1987) Why a diagram is (sometimes) worth ten thousand words. *Cognitive Science,* **11**, (1), 65-100.

Levie, W.H. and Lentz, R. (1982) Effects of text illustrations: a review of research. *ECTJ,* **30,** 4, 195-232.

Lopuck, L. (1996) *Designing Multimedia.* Peachpit Press.Berkley, CA.

Mayer, R. E. & Anderson, R.B. (1991) Animations need narrations: an experimental test of a dual-coding hypothesis. *Journal of Educational Psychology,* **83**, 484-490.

Norman, D. (1993) Cognition in the head and in the world. *Cognitive Science,* **17** (1), 1-6.

Reid, D.J. and Beveridge, M. (1986) Effects of text illustration on children's learning of a school science topic. *British Journal of Educational Psychology,* **56,** 294-306.

Reisberg, D. (1987) External Representations and the advantages of externalizing one's thoughts. *Proceedings of the IX Annual Conference of the Cognitive Science Society,* Seattle, W.A., Hillsdale, NJ., LEA

Rieber, L. P. (1989) The effects of computer-animated elaboration strategies and practice on factual and application learning in an elementary science lesson. *Journal of Computing Research,* **5** (4), 431-444.

Rogers, Y and Aldrich, F. (1995) In search of clickable Don Normans: learning through interacting with a CD-ROM 'book'. *Unpublished manuscript.*

Scaife, M. and Rogers, Y. (1996) External Cognition: How do graphical representations work? To appear in: *International Journal of Human-Computer Studies.*

Schwartz, D.L. (1995) The emergence of abstract representations in dyad problem solving. *Journal of the Learning Sciences*, 4, 321-354.

Shipley, C. (1996) *Anders CD-ROM Survey,* Andiron Press, Inc, Brookline, MA.

Soloway, E., Jackson, S. L., Klein, J., Quintana, C., Reed, J., Spitulnik, J., Stratford, S. J., Studer, S., Jul, S., Eng, J., & Scala, N. (1996) Learning Theory in Practice: Case Studies of Learner-Centered Design. *ACM CHI '96 Human Factors in Computer Systems, Proceedings,* Vancouver, B.C.

Smith, R. (1986) The Alternate Reality Kit: An animated environment for creating interactive simulations. *Proceedings of IEE Computer Society Workshop on Visual Languages.* 99-106.

Stenning, K. & Oberlander, J. (1995) A cognitive theory of graphical and linguistic reasoning: logic and implementation. *Cognitive Science*, **19**, 97-140.

Stenning, K. and Tobin, R. (1995) Assigning information to modalities: comparing graphical treatments of the syllogism. In: E. Ejerhed & S. Lindstrom, Eds. *Action, Language and Cognition.* Proceedings of International Conference on Dynamic Semantics, Umea: Sweden. September 1993. Kluwer.

Turkle, S. (1995) *Life on the Screen.* Simon and Schuster, New York.

Vera, A.H. and Simon, H.A.(1993) Situated Action: A Symbolic Interpretation. *Cognitive Science,* **17** (1), 7-48.

Woolf, B.P. and Hall, W. (1994) Multimedia Pedagogues: Interactive systems for teaching and learning. *Computer,* May 1995, 74-80.

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

Red coding: Articles and interactive demos available on the Web

Johnston, W., Nip. W. and Logan, C. (1995) "Whole Frog" Project.
<http://george.lbl.gov/ITG.hm.pg.docs/dissect/info.html)

Sims, R. (1995), Interactivity: A Forgotten Art?
<http://kvoneida.hbg.psu.edu/FACULTY/Farquhar/DLPE/ITFORUM/itfosims.htm>

Soloway, E. Jackson, S., Klein, J., Quintana, C., Reed, J., Spitulnik,J., Stratford, S. and Studer, S. (1995) Learning Theory in Practice: Case Studies of Learner-Centered Design.
<http://www-personal.umich.edu/~spit/Hi-C/DIS.html>

Cyan coding: On-line CD-ROM reviews

CD SourceLine
<http://www.omix.com/cgi-bin/cds/index.cgi>
Bruce, Linda (1995) Multimedia encylopaedias ("Material reproduced from The Edutainment Page", Copyright)
<http://www.edutainment.com.au/encycli.html>
misc.kids.computer (Collection of reviews found in newsgroup)
<http://qv3pluto.leidenuniv.nl/steve/reviews/welcome.htm>
Online CD-ROM magazine
<http://www.widearea.co.uk/cdrom/>
Elko, Steve and Anzovin, Lance (1996) Twin Peaks. "Look at the Current Releases"
<http://www.omix.com/cgi-bin/cds/cds.cgi?749248+articles/tp5.html>

Pink coding: CD-ROMs

BodyWorks 4.0 (1993-1995) Softkey International.
Earth Explora (1995) Enteractive Inc and Apple Computer, Inc. CA.
Encarta'96 (1996) Microsoft Corporation.
First Person (1995) Norman, D. Voyager.
<http://www.voyagerco.com/CD/ph/p.norman.html>
The Random House Kid's Encyclopaedia (1996) Knowledge Adventure.
Way Things Work (1995) Dorling Kindersley Publishing.
What's the Secret? (1995) 3M Learning Software.

