

Improving Children's Understanding of Formalisms through Interacting with Multimedia

Yvonne Rogers, Mike Scaife, Frances Aldrich & Sara Price
(*School of Cognitive and Computing Sciences, University of Sussex, Brighton, BN1 9QH, UK*)

Cognitive Science Research Paper 559
2003

Abstract. IMM (Interactive multimedia) has the potential to facilitate learning by providing new means of interacting with information, offering learners the ability to explore ideas and concepts that they find difficult to understand when represented in traditional media (e.g. diagrams and text in books). To know how to capitalise on this, however, requires understanding how IMM works. In this paper we describe our Cognitive Interactivity framework that outlines the benefits and properties of IMM. To support our assumptions we developed and empirically tested a software prototype for teaching children how to use foodweb diagrams to reason about the dynamic behaviour of ecosystems. Findings suggest that IMM can be developed to help children understand better how to use formalisms as computational aids with which to reason about a complex system.

Keywords: interactive multimedia, computational offloading, dynalinking, children, learning, diagrams, ecology, external cognition.

Introduction

Interactive multimedia (IMM) offers instructional designers an unprecedented opportunity to create rich interactive learning environments. Animations, diagrams, text, speech and video can be combined together to provide interactive simulations and other 'hands-on' interactivities that support learning in ways not possible with traditional media (e.g. books). In particular, IMM has the potential to provide new means of interacting with information, offering learners the ability to explore ideas and concepts in innovative ways. Much educational software has been developed to exploit this capability (e.g., see Hartley, 1994, for review). These include the provision of physics simulations, arithmetic games and hands-on experiments. A key question, however, is whether IMM per se does actually facilitate learning and if so, how.

Numerous evaluation studies investigating the effects of multimedia-based aids on learning show mixed findings, on some occasions demonstrating they have been more effective than traditional learning materials whilst on others, to be no better and sometimes worse (e.g. Narayanan and Hegarty 1998; Jones and Scaife, 1999; Pane, Corbett & John 1996). Furthermore, accounts for these effects have been localised and difficult to generalise from. Attempts to explain how interactivity can facilitate learning have also been limited (Kirsh, 1997; Sims, 1994). Little is known, therefore, about how to predict the conditions under which IMM is an effective aid to learning and moreover, what particular characteristics of multimedia promote these advantages (Scaife and Rogers, 1996).

In order to determine how IMM can be designed more systematically as effective learning environments it is important to have firstly, a theoretical framework, which can inform design decisions about how to combine and represent information using IMM, and secondly, a better understanding of the process of learning through IMM, to determine which kinds of learning activities will benefit most. Our research is concerned with developing such a theoretical framework and using it to inform the design of IMM for supporting topics which have been found to be difficult to learn when using only traditional media. In particular, we focus on subjects which are (i) complex, having a number of interdependencies and interrelationships, (ii) where the use of abstractions is necessary to understand and reason about them, and, (iii) where children find it difficult to learn how to use these to reason effectively about the underlying topic. An example of a difficult topic, which we are concerned with here, is ecological concepts. These are complex dynamic systems, that in order to reason about their behaviour requires the use of a number of formalisms. Previous research has shown that children have considerable difficulty understanding both how these kinds of dynamic systems work and the formalisms that are an inherent part of the subject area (e.g. Griffiths and Grant, 1985; Johnstone and Mahmoud, 1980; Scaife et al, 1997). Empirical studies suggest that a main problem is that learners find it hard to understand how to map between the different levels of abstraction provided by the formalisms and their existing everyday knowledge of the domain.

For the area of ecology, our claim is that IMM can be developed to enable learners to understand how to use the formalisms in the way intended, i.e. as computational tools with which to reason about the behaviour of a dynamic system. It can do this through showing how the various abstractions map onto the underlying domain. In particular, we believe that the dynamic representational properties of IMM provide much scope

for explicitly showing how such abstractions can be used as reasoning aids. Our claim about the cognitive benefits of IMM for supporting learning is derived from our theoretical framework called Cognitive Interactivity (Scaife and Rogers, 1996). Essentially, this framework describes the cognitive processes by which new information is integrated with existing knowledge and re-represented, in terms of the coordination of internal and external representations. The framework has been operationalised to characterize the various properties of external representations, enabling hypotheses to be made about which combinations of them will support effective learning (see section 2 for more details).

To test our assumptions about the benefits and properties of IMM for promoting learning we developed a software prototype for teaching children how to reason with an abstract formalism. Based on our cognitive interactivity framework, we designed a suite of IMM modules aimed at teaching children how to use foodweb diagrams to reason about the dynamic behaviour of ecosystems. An empirical study, where pairs of children used the software was then carried out. Findings from pre and post tests showed significant improvement, supporting our hypothesis that IMM can be developed to help children understand how to use abstractions as computational aids with which to reason about a system. This paper describes our approach to developing effective IMM. Firstly, we describe the specific kind of learning we are interested in facilitating, focusing on the importance of integrating knowledge through multiple representations. Secondly, we present our theoretical framework called Cognitive Interactivity, that outlines the cognitive benefits of interacting with external representations. Thirdly, we explain how we operationalised our framework in relation to our model of learning, and in particular, in terms of how we designed IMM for teaching children about ecosystems using foodweb formalisms. Finally, we describe our empirical study where we tested our assumptions about the benefits of cognitively-informed IMM for learning, in particular, its ability to promote reasoning skills.

The learning process: Integrating knowledge through interacting with multiple representations

Our approach to learning through using IMM resonates with the Learner Centred Design (LCD) framework (e.g. Soloway et al, 1996). A central claim from this perspective is that software can be an effective means of providing scaffolding for the learner, helping them to engage in activities that are normally beyond them. A central role of the software is to support the learner in engaging in these activities so that they

can master them. In so doing it can reduce the complexity of the activity. A further assumption behind this approach is that a learning environment should enable learners to interact with representations that are grounded in their previous experience (Jackson, et al, 1996). For example, the Model-It software, developed based on LCD principles, initially provides familiar objects for learners to build more abstract simulations of dynamic systems.

Likewise, we assume that learning is best supported through providing scaffolding, such as familiar representations that can be used as a basis from which to explore unfamiliar ones. A key research question this raises is what is the best way to coordinate different kinds of representations to support learning. A common strategy in classroom teaching is to get students to interact with and use multiple representations (e.g text, diagrams, pictures) when learning about a topic. For example, a bar chart and a pie graph may be shown together depicting the same mathematical information. A further assumption is that learning about a domain through multiple representations can engender different ideas whilst also helping the learner constrain their interpretations (see Ainsworth et al, 1997 for a review). Furthermore, it is often argued that the more appropriately different representations are integrated by the learner the more likely that ‘deeper’ understanding will occur (Kaput, 1989; Laurillard, 1993).

IMM offers an obvious platform for presenting such combined representational environments. Indeed, a number of computer-based systems have been built specifically to provide multiple representations of the same information. For example, Soloway et al (1996) have developed various modelling environments (including Model-It) which enable children to build models, such as stream ecosystems, using physical objects between which they define relationships. The system then shows these relationships as both textual and graphical representations. Hyperproof is another multimedia environment that was developed by Barwise and Etchemendy (1994) for teaching first order logic. This system also presents both text-based and graphical representations that are interlinked to help students work out proofs for logic problems: the rationale being that providing this combination of representations allows students to capitalise on the complementary properties of the different modalities, using them concurrently or switching between them.

As Underwood (1997, p. 4) observes, however, the general strategy of providing multiple representations to support more effective learning is not without its problems: “... is this a ‘good’ thing? Is this an additive process in which more

(representations) always mean better or more efficient (learning), or does the use of multiple representations place new learning demands on the child?" This concern is also addressed by Ainsworth, Bibby & Wood (1997) in their study of children using a computer-based system which employed mathematical and pictorial representations to teach numerical estimation skills. They observed that groups given mixed external representations (pictorial and mathematical) did not improve as strongly as those given multiple external representations of a single type. Much of the empirical research on the benefits of presenting information in redundant modes compared with using single media (e.g. spoken or written text with and without static diagrams or animations) has also been inconclusive (for a review, see Levie and Lentz, 1982). In particular, a number of studies have found that people find it difficult to integrate information presented in different modalities such as 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.

One of the reasons learners find it difficult to integrate multiple representations is that it requires additional work: they have to both interpret an individual representation in the domain and to translate between the different representations (cf. Cox and Brna, 1995). Furthermore, the mappings between the representations may not be obvious to the learner, making it confusing to switch between them. A key issue, therefore, when providing multiple representations for learning is to consider how to support this translation process so that it is more obvious and explicit to the learner. Another reason why multiple representations may be less effective than assumed is that they may simply be providing the same information at the same level but in different modalities, and thus providing no real means by which to obtain a deeper understanding. If the representations were to depict the same concept but at different levels of abstraction then perhaps this would lead to deeper understanding since 'deeper' levels of description were being provided. Another key issue, therefore, is whether providing multiple representations to convey different aspects of the same phenomenon may be a more effective strategy than simply representing the same information in different modalities. Our contention is that multiple representations are likely to be most beneficial when used to depict a concept at different levels of abstraction, for example, combining a concrete everyday simulation with an abstract formalism. Used in this way, each representation can provide a different perspective but also map onto each other, guiding the learner to reflect on the relationship between them. However, simply displaying representations at different levels of

abstraction will not by itself enable the learner to understand the relationship between them. What is also needed is a way of allowing the learner to actively explore the mappings. A further contention is that IMM can fulfil this role by providing the means by which multiple representations can be explicitly and dynamically linked with each other, together with the possibility of conveying mismatches and other conflicts between them, in ways that traditional media simply cannot. Before outlining how this can be achieved we first identify a problem domain, which learners currently have difficulty understanding.

The learning problem: reasoning with formalisms about a domain

An integral part of any science domain is the many kinds of formalisms, such as cycle diagrams and flow charts, that have evolved to allow predictions and inferences to be made about the interrelationships between elements and processes for a given concept. For example, webs are used to convey feeding relationships and energy flow in ecosystems; predictions can be made about what will happen to the ecosystem when it is perturbed (e.g. a population is removed) by appropriately reading the elements represented in the diagram. Yet, typically, children are not taught the diagram-reading skills necessary for reasoning with such formalisms (e.g. see Lowe, 1996; Cox, 1996), with the result that notations such as arrows linking the parts of a food web diagram are not properly understood. Consequently, children find it difficult to make correct predictions from formalisms such as foodweb diagrams. Instead, they tend to rely on their intuitions as to what different symbols mean and how to use them to make inferences, often wrongly since the conventions used in scientific diagrams can be counter-intuitive to their everyday assumptions.

Ecology is a domain that is replete with formalisms, representing a range of concepts at different levels of abstraction. These include food webs, pyramid of numbers and energy cycles. It is hardly surprising, therefore, to find that children have much difficulty using them, frequently misinterpreting them, ending up with a number of misconceptions about ecosystems (Johnstone and Mahmoud, 1980; Griffiths and Grant, 1985; Scaife et al, 1997; Rogers and Scaife, 1998). This state of affairs is particularly problematic in the UK, since the National Curriculum for education targets understanding of ecosystems to a level where students are required to use formalisms such as the foodweb diagram to make complex inferences about cause-effect relationships (e.g. if parameter X changes then what happens?) (see Figure 1). Instead, what tends to happen is that a chasm develops between the child's real world knowledge and the semantics of formalisms, whereby the child never really gets to

grip with understanding how to use the formalisms as computational tools with which to reason about the real world. Whilst, on the one hand, they have no problem understanding what the food web diagram represents, i.e. that within an ecosystem different organisms some but not others (e.g. in a pond fish eat tadpoles and tadpoles eat weed but that the tadpoles don't eat the fish or the weed eat the tadpoles), on the other they find it difficult to use the food web diagram, to reason about the ecosystem as a whole. Often they read or draw the diagram to represent incorrect relationships between organisms. Furthermore, they are often unable to read information from the diagram about which organisms die as a consequence of one of the other organisms being removed from the ecosystem.

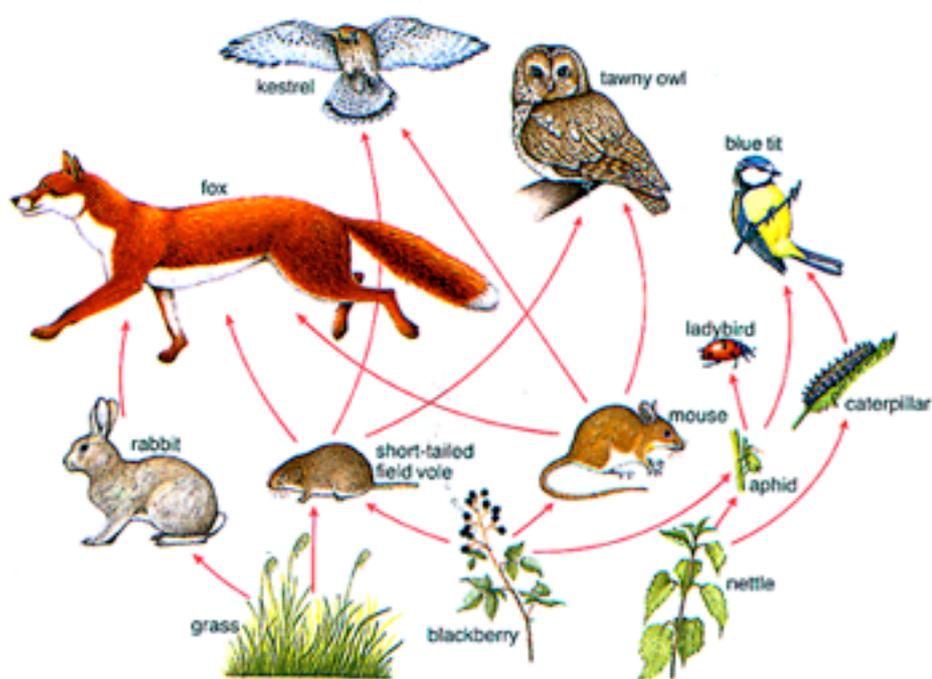


Figure 1. A typical food web diagram used in school textbooks to show who eats what in an ecosystem. Here one of the organisms has been shaded out in the diagram in order to pose the question, “The mice have died. What would happen to the rest of the ecosystem?”

The ability to do this kind of inferential reasoning – using the spatial layout to ‘read off’ the solution, by working through the chains of arrows between organisms – is a fundamental part of understanding ecological concepts (as opposed to simply memorising a food web diagram). Here, therefore, is a problem domain where we considered the learning process could be significantly improved through the use of IMM. In particular, the computational and representational functionality provided by IMM could be exploited to enable the child to learn to understand and reason more

effectively with the formalism of the foodweb about the behaviour of the ecosystem.

Theoretical framework – Cognitive Interactivity

The theoretical framework that we used to inform the design of our IMM for learning about ecosystems is ‘Cognitive Interactivity’ (Scaife and Rogers, 1996; Rogers and Scaife, 1998). This approach emphasises (i) the process by which new information is integrated with existing knowledge and then re-represented and (ii) the cognitive benefits and costs of particular forms of representation. The framework allows us to identify the properties of external representations in terms of their ‘computational offloading’. This refers to the extent to which different external representations reduce or increase the amount of cognitive effort required to understand or reason about what is being represented. High computational offloading is where much of the effort is offloaded onto the representation, requiring minimal effort on behalf of the learner for a given task. In contrast, low computational offloading is where much cognitive effort is required by the learner. In our analysis we have identified three main forms of computational offloading (Scaife and Rogers, 1996). These are:

- *re-representation* – This refers to how different external representations, that have the same abstract structure, make problem-solving easier or more difficult (see also Peterson, 1994, 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.
- *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 (see also Stenning and Tobin, 1995; Stenning and Oberlander, 1995).
- *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.

Together, these cognitive characterisations provide us with a basis from which (i) to predict the effects of combining different external representations and (ii) to specify the properties and trade-offs of combining different representations for different tasks. Hence, through exploiting different constraints different combinations of representations can be designed such that relevant elements and relations between

them can be made explicit for a given stage of the learning task.

Operationalising cognitive offloading in relation to the learning process

Computational offloading: task demands: At a general level, we operationalised computational offloading in terms of the amount of cognitive work the learner was expected to do at different stages of a learning task. For the first learning activity we decided that a high level of computational offloading was important whereby the learner is simply required just to interact with a simulation to discover things about what is being represented (e.g. feeding relationships of organisms in an ecosystem). In subsequent modules the level of computational offloading was generally decreased, requiring the learner to do increasingly more cognitive work as their understanding of the domain increased. For example, in a later module an empty template is presented which the learner has to complete by placing appropriate elements in the correct place. This requires them to make inferences about the ecosystem by interacting with the formalism – requiring a higher level of understanding.

Computational offloading: dynalinking A key form of temporal and spatial constraining which we investigated in this study was dynalinking. This is a specific property of IMM that static representations do not have – whereby multiple representations can co-vary with each other over time and space such that making changes to elements in one display are shown to co-vary in another kind of display. The control of this coupling is initiated through the learner; the computer system displays the consequences in another representation(s). For example clicking on an arrow in a food web diagram shows the feeding behaviour it represents in an adjacent dynamic simulation of a pond (see Figure 2). Dynalinking can also be used at higher levels of abstraction, such as to convey the knock-on effects within an eco-system when it is perturbed.

One of the main cognitive benefits of dynalinking is to allow relationships between elements of a complex concept(s) to be dynamically and explicitly displayed, together with the possibility of conveying mismatches and other conflicts between them. In relation to learning, our prediction is that it can firstly help learners integrate mappings between representations at different levels of abstraction, and secondly, support them in understanding better how to reason with the formalisms of a domain.

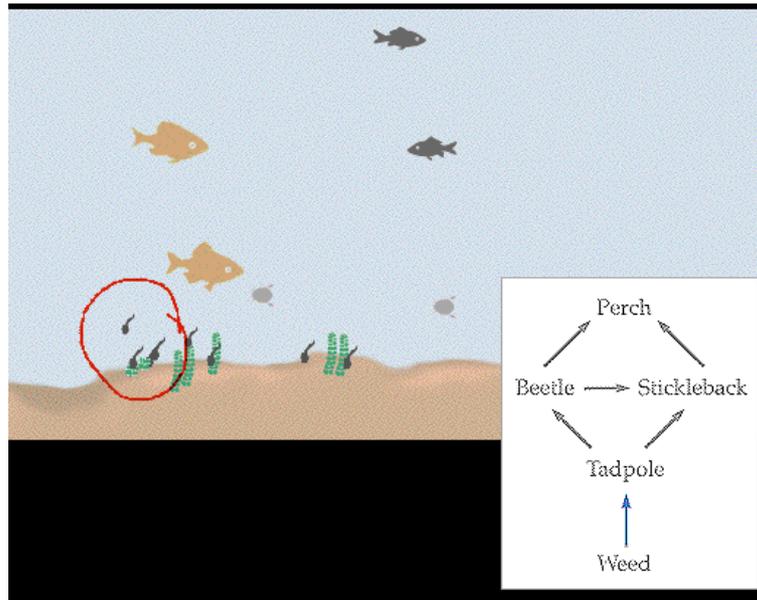


Figure 2. A snapshot taken from the PondWorld software prototype to demonstrate dynalinked representations. (The blue link in the food web diagram corresponds to the highlighted action of the organisms in the adjacent animation).

Designing the software prototype: Implementing our ideas about computational offloading

A software prototype, called PondWorld, was implemented, comprising a suite of interactive modules that depicted an ecosystem of a pond at varying levels of abstraction. Each module provided different kinds of interactivities, including exploring, constructing and manipulating abstractions for given scenarios. The software modules were developed to provide an appropriate level of scaffolding for the various stages of the learning task. Learners were required to complete the problem-solving tasks set in each module before moving onto the next one. Accordingly, the activities in each module were designed to vary in terms of the amount of cognitive effort required by the learners to accomplish them. Each module was also designed to increase in complexity - in terms of what was being represented and what problem-solving activities needed to be solved. An underlying pedagogical rationale was to allow learners to integrate the new knowledge presented in each module with what they had already learned from the previous modules and to be able to re-represent it at higher levels of abstraction. In sum, each module was operationalised in terms of:

- level of computational offloading
- form of multimedia interactivity
- problem-solving task
- learning process supported

Module 1: PondWorld Simulation

A concrete representation of a simple ecosystem was provided in the form of an animation of a pond with a small number of inhabitant species (see figure 3). The PondWorld animation showed fish predators eating water beetles, water beetles eating tadpoles and tadpoles consuming weeds. The learners interacted with the animation by clicking on the various organisms. When activated each organism tells the learner what it is and what it eats. The voices used were designed to vary in pitch, from low to high - a design idea based on a suggestion by the children during an informant design session (see Scaife et al, 1997). For example, the largest predator when clicked on, says in a deep voice: “I’m a perch, I eat beetles and tadpoles”. Conversely, the primary food source when clicked on, says in a high pitched squeaky voice, “I’m a weed. I make my own food.”

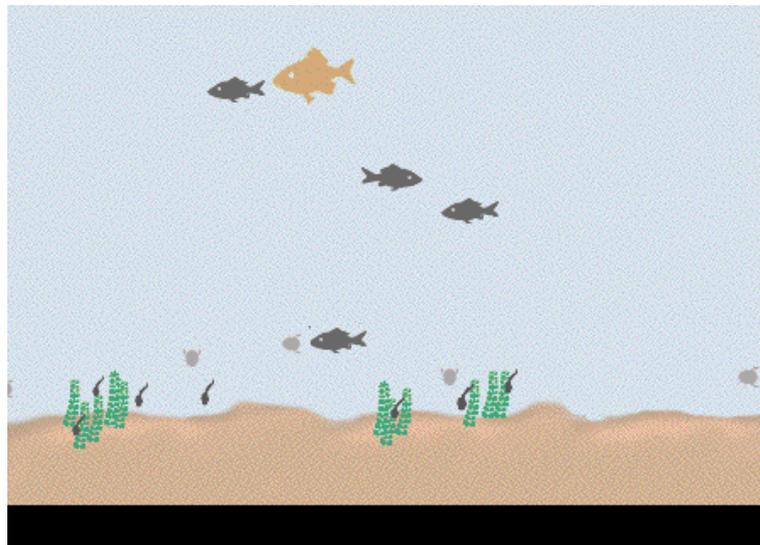


Figure 3. A snapshot of Pondworld Simulation, showing organisms eating each other. The organisms can be clicked on to discover what they eat (through the use of spoken narrative)

The level of computational offloading in this module is high: learners have only to interact with the animation by pointing, clicking and listening to the voices. The purpose of providing this kind of interactivity was to guide the attention of the learners to the different relationships within the ecosystem.

The problem-solving task: After interacting with this module a multiple choice quiz is presented on the screen to allowing the learner to immediately test their knowledge of the feeding relationships. The rationale for including this was to ensure that the learners had a correct understanding of the feeding relationships within PondWorld before moving onto the next module. The system asks a series of questions and lets them know immediately whether they are correct or not: For example, for the question “what does the stickleback eat?” the answers can be either correct (tadpole and beetle) partially correct (tadpole) or incorrect (weed). Providing a partially correct answer as an option gives the child the opportunity to reflect on why this is not quite correct and in so doing enable them to understand that organisms can eat one or more organism. The learning process supported in this module is obtaining factual knowledge; i.e. feeding relationships between a set of organisms in a given community.

Module 2 IntroWeb

In this module the learner is presented with two adjacent representations: a canonical food web diagram and a concrete simulation of it (see figure 4), the former being an abstraction of the latter. Narration is provided at the beginning to explain the relationship between the two forms of representation. The two representations are coupled using dynalinking: the organisms in the animation are designed to behave in relation to the abstract feeding relationships depicted in the food web. Here the level of computational offloading is still relatively high. The problem-solving task requires the learner to select different feeding relationships (as represented by the arrows) in the food web formalism and observe the outcomes of their action in the concrete animation. For example, clicking on the arrow link between the weed and tadpole in the foodweb diagram, results in the animation showing a token weed slowly being chomped by the tadpole. The rationale behind this module is to draw the learner’s attention to the mapping between the formalism (i.e. the food web diagram) and the concrete animation of the same underlying referent (i.e PondWorld) through dynalinking. After familiarising themselves with the task learners were asked to make predictions about what would correspondingly happen in the pond before clicking on an arrow.

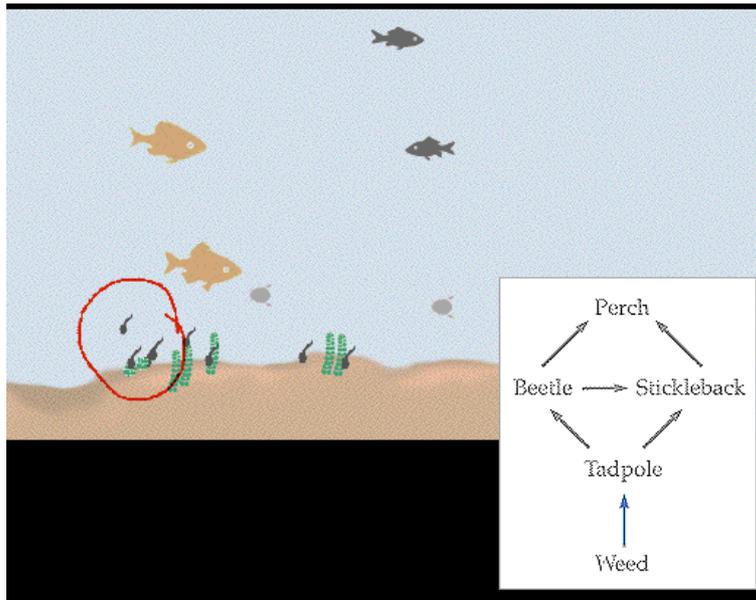


Figure 4 A snapshot of IntroWeb, showing dynalinking of the Pondworld simulation with a foodweb diagram

The learning process supported in this module consists of two interlinked key aspects: (i) understanding what the canonical forms used in foodweb diagrams represent and (ii) learning the mapping between these and the organisms and implicit processes represented in the concrete animation (i.e. the PondWorld ecosystem). To achieve (i) requires understanding that the organism at the head of an arrow is eaten by the organism at the tail of the arrow, i.e A is eaten by B (see Figure 5). A typical misconception that children make is to read the diagram as A eating B, following their intuition about the directionality of a process represented by an arrow. In the context of the foodweb diagram this is an incorrect interpretation. A further aim with this module, therefore, was to expose this misconception to learners through the explicit linking of the diagram with the concrete simulation they were already familiar with.



Figure 5 Canonical form used in food web formalisms where the arrow represents A is eaten by B

Module 3: LinkWeb

In this module the learner is presented with a more complex ecosystem. More organisms have been added to PondWorld which are depicted both in the simulation and a template of the food web diagram. The initial learning activity is to enable the learner to recognize that both the animation and the formalism have changed and for them to map these onto each other. The change is made explicit through narration, and mapping can be achieved by clicking on new organisms, as in Module 1.

The problem-solving task: The diagram is presented as a template for the learner to fill in by placing arrows to indicate the feeding relationships between the various organisms (see figure 6). The learner's task is explained through narration at the beginning of the module. There are 8 links in the diagram that have to be completed. This is done by clicking on the organisms in the correct order (e.g. slime eaten by snail) on the diagram. Feedback is displayed in the form of coloured arrows, which appear when the correct feeding relationships have been linked. The learner can verify who eats what by clicking on the organisms in the adjacent simulation to hear them speak. Again the use of dynalinking was included to encourage the learner to make explicit links between the different levels of abstraction.

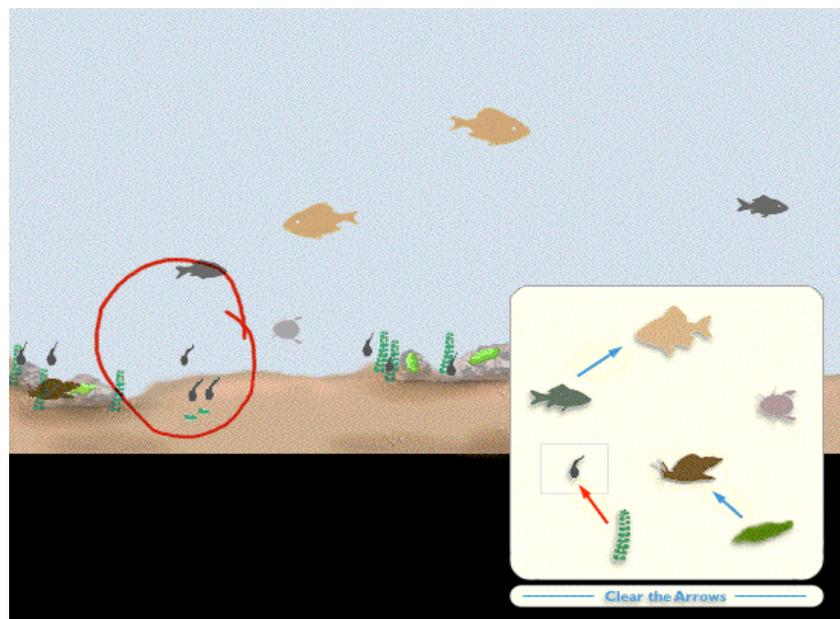


Figure 6. Snapshot of LinkWeb showing partially constructed foodweb diagram dynalinked with PondWorld simulation

In some foodweb diagrams the size and shape of the tokens can often give clues as to the feeding relationships: smaller ones are eaten by larger ones and those positioned below are eaten by those above. To prevent the children from simply using these

dimensions when reasoning with the diagram we designed the tokens for the organisms to be roughly the same size and also for both horizontal and vertical arrows to be included. This way they would have to make inferences about the ecosystem on the basis of understanding what the arrows and tokens represented.

In this module the level of computational offloading has been decreased in so far as the learner is required to make a number of new inferences about the domain and the formalism: (i) that ecosystems can increase in complexity and that the food web diagram is designed to show this at a higher level of abstraction through the use of directed arrows between the organisms represented in the diagram and (ii) to partially construct a food web diagram by working out the correct links between the organisms at a higher level of abstraction. The learning process supported in this module again consists of: (i) understanding what the canonical forms used in foodweb diagrams represent and (ii) learning the mapping between these and the organisms and implicit processes represented in the concrete animation (i.e. the PondWorld ecosystem). Compared with the previous module, however, the learners are required now to construct their own foodweb diagrams with the new assortment of organisms in the ecosystem, i.e. to generalise their learning to a new situation.

Module 4: EraserWeb

The fourth module of PondWorld, called EraserWeb, was designed to show the same two interlinked representations as in LinkWeb. This time, however, the learner was required to infer what would happen to the ecosystem when it is perturbed, i.e. when one of the species is removed. The objective here is to get the learner to reason about the ecosystem (i.e. what the consequences will be for the other organisms) by reading off and interacting with the foodweb diagram (see figure 7). To prevent the children from simply recalling the spatial positioning of the tokens from the previous module the spatial ordering of the species was switched around.

An example of such a perturbation is demonstrated initially to the learner through a narrated animation of what would happen to the ecosystem when the tadpoles are removed. Crosses are placed serially on the organisms in the diagram to indicate the knock-on effects throughout the ecosystem as a consequence of the tadpoles being depleted. At the same time their concrete counterparts in the adjacent simulation are removed from the pond – the module again emphasising the dynalinking between the two forms of representation. The demonstration is then repeated to emphasize specifically the order in which the organisms will die off.

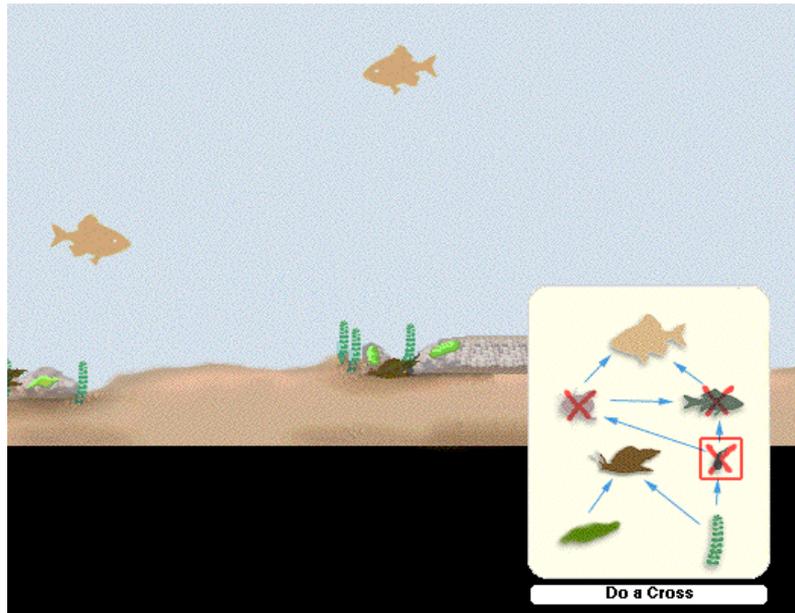


Figure 7. Snapshot of EraserWeb module where the tadpole has been removed in both the diagram and the simulation. The problem-solving task is to work out which other species would die off as a consequence, by placing crosses on the token representing those species in the diagram.

Problem-solving activity: The learner is then presented with two problems to solve by themselves, by working out which other organisms will die off as a consequence of one of the species being removed from PondWorld. One of the problems is relatively easy to solve (when the weed is removed and all other organisms die off) whilst the other requires a higher degree of inferencing (the water beetle is removed and only certain ones will die off). To solve both problems the learner is required to drag and drop crosses from an adjacent palette onto the organisms they think will die in the food web diagram. Learners were also required to verbally make predictions of what would happen to certain organisms before placing their crosses, or were required to explain why they placed crosses in the order they chose. Hence the level of computational offloading is low: even though the diagram provides the means by which to reason about the ecosystem the learner needs to know how to use it and to remember the routes they have followed throughout the Web to ensure they have explored all possible knock on effects on the other populations.

Learning process: In this final module, emphasis is placed on getting the learners to make inferences between the two interlinked representations and to understand that the diagram is an abstraction of the other that has more computational power allowing the solution to be systematically ‘read off’.

To summarize, Table 1 illustrates the main differences between the PondWorld software modules with respect to (i) the overall level of computational offloading in terms of cognitive effort required by the learner to complete it (ii) the main form of multimedia interactivity, (iii) the problem-solving activity, and (iv) the learning process involved in each.

Module	Computational offloading	Form of MM interactivity	Problem Solving Activity	Learning Process
1. PondWorld Simulation	High	Click and tell animation	Learning of feeding relationships in ecosystem	Factual knowledge: feeding relationships
2. IntroWeb	Medium	Dynalinking between animation and formalism	Make links between animation and formalism	Canonical forms in foodweb & mapping between abstractions
3. LinkWeb	Medium	Dynalinking between animation and formalism	Complete partially constructed formalism	Canonical forms in formalism & mapping between abstractions
4. EraserWeb	Low	Dynalinking with facility for annotating diagram	Use formalism to reason about ecosystem behaviour	Temporal effects of extinction of species

Table 1: A summary table of the different modules that were designed for PondWorld characterised in terms of computational offloading, multimedia interactivity, learner activity and learning process.

The empirical study - testing our hypotheses about the benefits of computational offloading

To determine whether our software was effective for supporting learning we carried out a quantitative study, using pre and post tests. We also performed a qualitative analysis on video recordings of children's behaviour whilst interacting with the software to examine in more detail the learning process that occurred. Given the design aims of the software we anticipated that:

- (i) The children should demonstrate an increasing ability to answer more complex questions about inter-relations between species as they progress through the modules.
- (ii) Experience with the software should result in a generalisable understanding of the abstract formalism of the foodweb, such that the child will be able to reason better about possible changes to the ecosystem purely on the basis of a diagram.

Method

Fourteen pairs of children, from local schools, worked through the PondWorld

modules. There were four male pairs, six female and four mixed, ranging in age from 9 years 3 months to 10 years 2 months (mean 9 years 10 months). They had covered some aspects of ecosystems in classwork, as part of the Keystage 2 science UK National Curriculum. This involved an introduction to simple food chains and discussion of basic distinctions between ‘plants’ and ‘consumers’ but not working with foodweb representations.

Children worked on PondWorld in pairs (pairs being chosen by class teachers on the basis of their ability to work harmoniously together), the reasons being that: (i) children of this age frequently work in this way in IT-based class lessons; (ii) pair work allows discourse between children, facilitating the collection of verbal protocols and (iii) joint work may encourage the developments of insight over individual work, e.g. Doise & Mugny (1998). Each pair began with an introduction to the aims of the software. The exercise was described in general terms as being concerned with our developing software to help “teach children about food chains” and their participation would help with the development process. Then each module was introduced, in the order and method described above.

The children were allowed to interact with the software with a minimum of intervention by the adult experimenter who sat separately but close by. Interventions were made whenever the children asked for help or when a pair apparently became stuck on the module, such as persisting with the ‘wrong’ method of clicking on pairs of species in the LinkWeb module. Such interventions were infrequent but were thought necessary to support the children’s’ progress through the software suite and to better mimic the possible use of such software in a teacher-led classroom context. All interventions, however, were as neutral as possible with respect to supplying children with any ‘correct’ answer to the tasks posed by the software. The time spent on each module varied with each pair, the total time ranging from 25 to 45 minutes. All pairs finished the entire series of modules and all sessions were videotaped. Children were given a pre- and a post-test as described below.

Pre- and post-tests to assess learning

The child’s ability to ‘read’ a foodweb diagram is, as we have indicated, a function of two things: the knowledge of what the links - arrows - mean, and a consequent ability to use these links to reason about relationships between species at some distance from each other in the web. However, when confronted with a foodweb containing familiar plants and animals, the child can use world knowledge about ‘who eats whom’ to identify relationships between adjacent species in the web. For example children

know that tadpoles eat weed, and not vice-versa, and can ‘read’ this from the food web regardless of their understanding of what the linking arrow might mean.

In designing a test for foodweb diagram understanding we, therefore, devised something that would help to factor out the possible influence of world knowledge of species behaviour. Initially we produced a foodweb that used pictures of imaginary animals at the diagram nodes (animals morphed from others or entirely made-up in a computer graphics programme). However many children, ingeniously, made inferences about likely feeding habits on the basis of cues such as whether the imaginary creatures had sharp beaks. Thus we moved to using an abstract foodweb diagram, shown in Figure 8. This was presented to the pairs as an A4-sized diagram, with blank pieces of paper at the diagram nodes and arrows connecting them as shown in the figure. To prevent the task being too abstract for the children, they were told that under each blank piece of paper there was a picture of a plant or animal. The diagram’s shape resembled that of the one that is presented in the last PondWorld module but with a partial reversal of positions of the organisms lower in the feeding chains. This was to reduce the possibility of rote transfer in the post-test but to maintain a food web of sufficient recognisability and similarity to ones they had experienced.

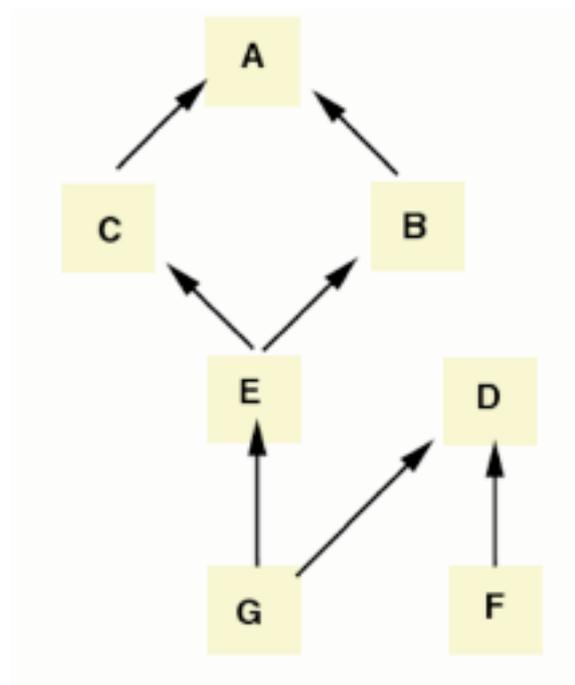


Figure 8. The abstract food web diagram, tests with characters A to G representing species, used in the pre and post tests

Each pair was presented with the foodweb test immediately before PondWorld and

again immediately afterwards. After explaining that the diagram represented a foodweb, the children were asked a series of questions to estimate their understanding of the formalism. The questions asked about the diagram were of three types, with three questions in each section as follows.

(a) who eats what?

This was a simple test of the children's ability to correctly read the arrows in the diagram. The questions were: 'what does D eat?', 'what does G eat?', 'what does C eat?'. The species vary in that C is likely to be a secondary consumer (carnivore), D a primary consumer (herbivore) and G a producer (plant). These questions require the children to use the arrows and the target species' position in the food web to identify what it eats. For G the correct answer is that it makes its own food (or doesn't eat any other organism).

(b) the effect of species deletion on the foodweb

Here the question was 'what will happen if we take away all the X', where X was E, F and B'. Correct answers require that children identify the knock-on effects of removal at a simple level. In this case removal of B will affect A but still leave A with an alternative food source; removal of F will similarly deprive D but leave it with G to feed on; removal of E, however, will ultimately result in the extinction of A, B and C.

(c) identifying higher-order feeding categories

This was aimed at the child's abilities to identify possible trophic categories within the foodweb, based on the canonical placement of species at different levels in the diagram. Thus plants (producers) are typically placed at the lowest level, primary consumers (herbivores) at intermediate levels and secondary consumers (carnivores) at the highest level. The questions were: 'what produces its own food?' (G and F), 'what eats only plants' (E and D) and 'what eats only meat?' (A, B and C).

In all cases the children were tested as a pair, with each child allowed to contribute to answers as they wished. Again this was thought desirable because of the importance of pair interaction and because experience with individual testing had revealed a far lower readiness to interact with the experimenter in completing the task. Where contradictory answers were given by the two members, such as different identifications of prey, a conservative scoring procedure was used and wrong answers were counted against correct ones. In point of fact of the fourteen pairs only four gave contradictory answers on any of the questions. For the vast majority of time either one member acted as a spokesperson or the children conferred before deciding on their answer. No feedback was given during either pre- or post-test as to whether answers were correct or not, the experimenter only offering general encouragement.

Quantitative analysis of pre- and post-test

The data from the pre- and post-test on the foodweb understanding task are presented in Table 2. A repeated measures ANOVA was performed on the scores from the pre- and post-tests and revealed a significant change in overall performance ($f = 13.46$; $df = 1$; $p < 0.003$) with no significant difference in performance between the three question categories ($f = .408$; $df = 2$; $p < 0.674$). Overall eleven pairs from fourteen (79%) showed some improvement on combined scores (mean of 6.5 correct responses over the pre-test) and three (21%) showed none.

	Question group 1	Question group 2	Question group 3
Pre-test	1.79 $sd = 1.05$	1.57 $sd = 1.45$	2.14 $sd = 2.32$
Post test	3.71 $sd = 2.89$	3.86 $sd = 2.41$	4.14 $sd = 3.35$

Table 2 Anova table for performance scores on pre and post tests.

How did children who started from different points on initial performance benefit from exposure to PondWorld? There was a strong correlation between initial and final performance for total scores on the foodweb test ($r = .66$, $p = .009$). However, underlying this, the picture was of a divided sample, with a continued poor performance of the initial four lowest scorers, the largest gains coming in the initial middle-scoring section of the remainder.

Qualitative analysis of learning and reasoning processes

The overall scores presented above show strong evidence of improvement in the post-test, considered across the sample as a whole. In this section we want to look in more detail at aspects of the behaviour of pairs whilst interacting with PondWorld and in the post-test to elucidate the learning process as a function of the provision of dynalinking and IMM.

Interacting with PondWorld: the ‘aha’ learning experience

As expected, one major reason why children made errors in the pre-tests and during the sessions when using the software arose from their misunderstanding of what the canonical links - the arrows in the diagrams - mean. As we described before, a typical misconception that children have is to read the arrow from A to B transitively, as A eats B, following their intuition about the directionality of a process represented by an arrow (see Figure 5). For example, in a classroom situation, one child from a pair shown a simple three item food chain gave the correct reading, but the other pupil said: "yeah, but the arrows are pointing *that* way, so the carrot eats the rabbit and the

rabbit eats the fox". This illustrates the clash between diagram convention and intuitive understanding. In the pre-test similar inaccuracies were common.

One key piece of evidence that the software was able to overcome their commonsense understanding of how to read arrows with the scientific convention of how they are used in food web formalisms, is that when talking through their accounts of what is happening in the food web, the other child might challenge them or they might observe the correct behaviour in the PondWorld animation. This seemed to have the effect of making them realise their misconception and helping them to correct it – a kind of ‘a-ha’ experience. Hence, we believe that pupils improve, in all the modules, due to two effects - the dialogue within the pair (as when child A corrects child B; some examples are given below) and the direct tutorial effects of the program. Observations from interacting with the last three modules indicate this:

Introweb

The children worked through this module easily, and clicked on all the diagram links to see the feeding relationships in the animation. Most pairs had several goes with each link and seemed to greatly enjoy the effects. Initially many pairs showed an overriding tendency to read the meaning of the arrow incorrectly with many instances of false predictions such as: "the beetle eats the stickleback, the stickleback eats the perch, the tadpole eats the beetle and the weed eats the tadpole". In some cases the clash between real-world knowledge and the arrow convention meant that pupils would not accept that the diagram had meaningful semantics, as for the arrow which points from the weed to the tadpole - "Nothing will happen because the weed makes its own food". However, by the end of the module, the majority of pairs were able to read the arrows correctly having observed the match between link and animation.

The tutorial effects of the program in this module can be seen in this example: "Nothing will happen because the weed makes its own food" (initial prediction) followed by - "Oh No! The tadpole's eaten the weed" (as a result of seeing the animation). One pair showed insight from the Introweb animation that generalised to a previously misunderstood interaction with a simple three item food chain seen in class before PondWorld: "Oh, so that's what that thing..... the arrows go *up* so the leaf gets eaten by the caterpillar and the caterpillar gets eaten by the bird". This seems a particularly pertinent example. The pupil had understood the arrows in a way that makes the diagram now make sense to him. This is expressed as

1. The arrows go up. So:
2. A "gets eaten by" B.

LinkWeb

This module is more difficult to learn than the previous one. Here, the pairs have to focus on the arrow direction in that they have to initially select the species eaten and then the species that does the eating. In some instances, there appears to be a problem of transferring learning from one module to the next. Despite the ‘a-ha’ learning experience that appeared to have taken place in the IntroWeb, some pairs still seemed to expect the arrows to go the other way round in LinkWeb – they clicked first on the feeder, then the food. This occurred despite pupils making correct predictions orally about what they expected the relationship to be, suggesting a strong persistence of common-sense understanding. For example, the following interaction took place between C and M while placing arrows:

C: I want to do ‘the snail eats the slime’.

C clicks on the snail then the slime.

M: No, you do that (*points to slime*) then you do that (*points to snail*).

Here C knows that the snail eats the slime, but attempts to put the arrow the other way around. Again, it seemed that the conventions of the diagram are at odds with common-sense perceptions. The difficulty with correct clicking order is further demonstrated by J and A, who after having correctly made three arrows then need reinforcement to achieve the task by asking:

You click on the food first do you?

However, the beneficial effect of dynalinking enabling learners to make explicit links is also evident. After having constructed an arrow by clicking on the slime and the snail:

M: Oh, look! You can see he’s trying to eat all the slime (*pointing to the snail in the animation*).

EraserWeb

The majority of pairs (ten of fourteen) completed this section with reasonable confidence, and good predictions of the order in which crosses should be applied i.e. that the creatures would successively die out because there was no food for them. The crucial observation here is that several of the pairs, began to overtly articulate how to use the food web diagram to determine what would happen to the PondWorld when different organisms were removed from it. This contrasts sharply with their inability to use the diagram in the pre-test to work out what would happen when various organisms were removed. As an example, in the pre- test, pair R and D thought something would happen but could not work out what that would be. When asked later in an equivalent question, whilst interacting with PondWorld, they were

able this time to use the interactive diagram to work out what would happen:

R: um...the tadpole eats the weeds

D: because it doesn't have the weed the tadpole will die and then the beetle now because it doesn't have a tadpole and then it would be ...um...

R: it can't be any more because the snail has the slime and the slime makes it own food

Another example where the combination of the pupils talking with each other and interacting with the modules facilitated reasoning happens between pair C and M. Here they discuss the implications of an organism eating more than one other organism. The system presents the scenario: the weeds have gone from the pond.

G replies: That (*pointing to the tadpole*) would die first because that eats the weed. And then one of those two would die (*points to beetle and stickleback*).

T: but that (*pointing to the snail*) eats weed as well.

G: That eats weed as well, so Yeah, but that eats slime as well, so that can survive on slime.

From these and other video extracts, therefore, it is clear that most of the pairs are working with the diagram by following the arrows, node by node, determining whether the tokens for each of the species would die off as a consequence of the token weed being removed - indicating that they are using the diagram in the way it is supposed to be used.

Insight revealed on the post-test: evidence of being able to reason more effectively after using software

The pattern of reasoning on the post-test was our benchmark for learning from the PondWorld software. We gave above an example of one pair's behaviour when working with the diagram by following the arrows, determining which species would die off as a consequence of another being removed. When subsequently faced with the same problem in the post-test, i.e. with a static diagram, the pair again showed how they had become competent in using the diagram effectively to make the necessary inferences. For example, when asked what would happen if the species represented by token F was removed, the pair had the following discussion:

Q: What would happen if we took away all of F?

D: the D would die because it doesn't have any food

R: I don't think it does. If F dies the D won't die because it's got the G still to have something to eat.

D: Oh yes

Clearly, the two are drawing inferences about the behaviour of the ecosystem on the basis of their prior knowledge about the representation of feeding relationship rules and working-out the knock-on effects, using information from the directionality of the arrows connecting the species tokens. We also saw how one of the pair makes an initial partially correct inference, which spurs the other one to challenge this and come up with the full answer. Interestingly the roles of the children change for the next question, where this time child R challenges child D's initial answer:

Q: What would happen if we took away all of B?

R: Well the A...

D: No, no, 'cause look the A would have the C so the A would be all right. And the E would die because it doesn't have any...

R: It won't die.

D: Yeah, it won't die.

These findings suggest, therefore, that the majority of children were able to generalise their understanding from using the specific foodweb diagram in PondWorld to a more abstract example and had become adept at using the formalism to make the necessary inference.

Children who did not perform well in both pre and post tests

We were interested in why a few of the pairs did not perform well in the pre-test and showed little if any improvement in the post tests. On looking at the video data, it appears that they also had difficulty in interacting with the software modules. For example, one pair, A and H, showed difficulty understanding arrow meaning on the abstract diagram despite demonstrating a clear understanding of feeding relationships in the concrete simulation. This pair showed an overriding tendency to interpret the arrow to mean "eats" in the Intronweb, and LinkWeb modules. Here more continuous direction was needed about how to make the arrows, as there was prevailing inclination to click first on the eater and then on the food.

Interestingly, the children who had most difficulty with reasoning in the post tests showed lack of concentration generally and an inability to focus on the important aspects of the software, instead showing interest in inappropriate aspects, such as whether there were sharks, or what the sand would say, and paid more attention to superficial aspects of the software, such as the graphics which were found to be "cool", and what kind of sound the computer would make if they did something

wrong. These children had lower concentration levels than many of the other children, reducing the amount of information they could absorb. They required more direction, and more specific feedback about their actions, thoughts and answers from the researcher than other pupils, suggesting that for some learners more specific direction and feedback may be important in enabling transfer of concrete to abstract concepts. Thus, for these pupils type and amount of feedback may be important in determining the level of understanding that they reach.

Discussion

We developed the PondWorld software modules in an attempt to use dynalinking to 'bring to life' the conventions and semantics inherent in a particular class of diagrams. This study suggested that the software enabled most children to improve their ability to 'read' such diagrams. It seems that the interactivity offered in each module facilitated learning within that section, and by the end of each module, most pupils are able to understand what the diagram was showing in relation to the processes occurring in the pond. However, the persistence of reading arrows from a common-sense point of view rather than the correct scientific way was also noticeable. Even when a pair had successfully completed one module they could apparently revert to a wrong reading at the start of the next. Why might this be so?

The first possibility is that the learning that occurs within modules is rather specific to each. This may not be wholly surprising in retrospect since the interface varied in each case, with different actions required (click link, place cross etc) as indications of having grasped the 'point' of the formalism. Another factor which is relevant here is that the language used to describe the feeding relationships (X eats Y) and the associated questions (e.g who eats what?) are transitive in form - they tend to suggest a directionality which may, for the child, be at odds with the correct, but intransitive, reading of the arrow 'Y is eaten by X'. For some children this linguistic factor may have contributed to the problem, even though it entirely matches the usual language of class and textbook presentation (the reason why we used it). It is, however, unlikely to be the whole explanation since work with static food chain diagrams (Scaife, in prep) has shown that even an intransitive wording does not guarantee greatly improved performance.

Another observation that requires attention is that PondWorld was not universally successful in increasing learning, some pairs showing no learning gains. Why was this so if the IMM was as effective as we thought? A number of reasons may explain this

variability. In the first place it fits with previous research with multimedia – that it is generally most effective for ‘good’ learners and that strong individual differences for its effectiveness exist which may relate to cognitive style differences (e.g. Dillon & Gabbard, 1999). Secondly we have to realise that the learning target set here, the comprehension of foodwebs, is well above the curricular target for this age level (some three years or so) and that full comprehension may rely on further cognitive developments and/or learning, such as greater experience with other diagram forms that require interpretation of causal relationships. Thirdly, and most importantly, the actual task of ‘full’ understanding will, as we have flagged, require an integration of concepts (e.g. arrow notation, causality, consumer/producer distinctions) that may require repeated exposures to differing content domains (such as other ecosystems). In this respect it is worth recalling that foodwebs also work at another more abstract level to convey information about energy flow (hence the arrow direction), in addition to feeding behaviours, and that this a concept about which children of this age have little understanding.

This leaves us with the question of understanding the mechanisms for the learning that did occur, both by the end of the software and that persisted to the post-test. Above we mentioned two likely factors - the interaction between children and the form of the dynalinked animations. Clearly, too, increased motivation due to the fun factor of the software is important, and voiced clearly by the words and behaviour of the children as in "it's much more easy when we've been on the computer. It gives you a few ideas." However, while the software was undoubtedly motivating it seems unlikely that this is a sole explanation. Children are often very motivated by school trips to ecosystems but, according to teachers, still have difficulties relating to diagrammatic representations of those systems after the visit.

With PondWorld, by contrast, we hoped that the interactivity provided by the program would be effective in allowing students to go ‘beyond the information given’ (Bruner 1973) to generate hypotheses to fill in the gaps in their understanding. We wanted to support incremental learning as well as the kind of ‘aha’ experience cited above, where the linked animation clearly triggers insight. In gradually reducing the degree of computational offloading (c.f. Table 1) we were generally able to move the children to a more persistent ability to map, read and construct a diagram. The discourse of the pairs, as we have exemplified above, often gave evidence of an increasing ability to construct a chain of reasoning and experimentation, prompted by what the children were seeing. They were able to hypothesize and test with immediate visual feedback in ways not possible with books or, necessarily, in real-world

ecosystems. The evidence is quite clear that there was, modulo individual variation, an increased ability to handle abstraction and to use the abstraction to make inferences about the represented domain.

Summary

In terms of interactivity we have argued that the concrete dynamic presentation of an abstract concept allows pupils to match the processes they see occurring in the pond with the abstract diagram alongside. Thus, the ‘explicitness and visibility’ of the interactive multimedia direct attention to the important aspects in each software module. It offers control for the learner to the degree that the learner can choose their own pace to progress through the program, and reinforce aspects that they feel least familiar with. For example, they can click on creatures to get information as frequently as required and in the order that suits the learner. Feedback is offered through demonstration of what will happen when the learner, for example, clicks on a particular arrow, or puts a cross on a particular creature.

Our findings, therefore, suggest that interactive multimedia can play a powerful role in facilitating learning, through the implementation of different forms of computational offloading. The particular form we investigated in our study was dynalinking, where we showed how it can act as a learning aid to help pupils map their familiar concrete knowledge with unfamiliar abstractions of this knowledge, represented as formalisms. Hence, this supports one of our main cognitive assumptions about the value of IMM – that it can allow learners to explore the mappings between multiple representations, through conveying different aspects of the same phenomena that are explicitly and dynamically linked with each other. In particular, the combination of immediate visual and auditory feedback, that co-varied across different representations of the same concept at varying levels of abstraction, helped pupils to examine their incorrect common-sense based perceptions of how to read the foodweb diagrams. They also helped them learn how to use the foodweb formalism as a computational tool with which to reason about the dynamic behaviour of the underlying system.

Acknowledgements

The authors gratefully acknowledge support from the ESRC Cognitive Engineering Programme, project L127 30 1001 66. Many thanks to Matt Davies for his outstanding design and Director skills. We also thank schools in East Sussex, UK. Demos of the software prototypes can be interacted with at www.susx.ac.uk/ECOi

References

- Ainsworth, S., Bibby, P. & Wood, D. (1997) Information Technology and Multiple Representations: new opportunities – new problems. *Journal of Information Technology for Teacher Education*, 6(1), 93-105.
- Barwise J. & Etchemendy, J. (1994) *Hyperproof*. CSLI Lecture Notes. Chicago: Chicago University Press.
- Bruner, J.S. & Anglin, J.M. (1973) *Beyond the information given: studies in the psychology of knowing*. Allen & Unwin.
- Cox, R. (1996) Analytical reasoning with multiple external representations. PhD Thesis. University of Edinburgh. Edinburgh, UK.
- 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), 239-302.
- Dillon, A and Gabbard, R (1999). Hypermedia as an educational technology: A review of the quantitative research literature on learner comprehension, control, and style. *Review of Educational Research*, 68 (3), 322-349.
- Doise, W., Mugny, G. & Perez, J. (1998) The social construction of knowledge, In U. Flick (ed) *The Psychology of the Social*. NY USA Cambridge University Press.
- Griffiths, A.K. and Grant, B.A. C. (1985) High School student's understanding of food webs: identification of a learning hierarchy and related misconceptions. *Journal of Research in Science Teaching*, 22, 421-436.
- Hartley, J. (1994) Designing instructional text for older readers. A literature review. *British Journal of Educational Technology*, 25(3), 172-188.
- Jackson, S. L., Stratford, S. J., Krajcik, J. S., Soloway, E. (1996) Making Dynamics Modeling Accessible to Pre-College Science Students. *Interactive Learning Environments*.
- Johnstone, A.H. and Mahmoud, N.A. (1980) Isolating topics of high perceived difficulty in school biology. *Journal of Biological Education*, 14, 163-166
- Jones, S and Scaife, M. 1999; Diagram representation: A comparison of animated and static formats. *Proceedings of ED-MEDIA 99* vol 1, 622-627
- 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*. Hillsdale, NJ: LEA.
- Kirsh, D. 1997 Interactivity and multimedia interfaces. *Instructional Science*, 25(2), 79-96.
- Laurillard, D. (1993) *Rethinking University Teaching*. London: Routledge.
- Levie, W.H. and Lentz, R. (1982) Effects of text illustrations: a review of research.

ECTJ, **30**, 4, 195-232.

Lowe, R. (1996) Background knowledge and the construction of a situational representation from a diagram. *European Journal of Psychology of Education*, 11(4), 377-397.

Narayanan, N.H. & Hegarty, M. (1998) On designing comprehensible interactive hypermedia manuals. *International journal of Human-Computer Studies*, 48(2), 267-301.

Pane, J.F., Corbett, A.T. & John, B. (1996) Assessing Dynamics in Computer-Based Instruction. *CHI'96 Electronic Proceedings*.

http://www1.acm.org:82/sigchi/chi96/proceedings/papers/Pane/jfp_txt.htm

Peterson, D. (1994) Re-representation and emergent information in three cases of problem solving. *Artificial Intelligence and Creativity*, 81-92.

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 & Scaife, M. (1998) How can Interactive Multimedia Facilitate Learning? *Proceedings Of First International Workshop On Intelligence And Multimodality In Multimedia Interfaces*. AAAI Press, Menlo Park, CA.

Scaife, M. (in prep) Children's understanding of food chain and food web diagrams.

Scaife, M. & Rogers, Y. (1996) External Cognition: How Do Graphical Representations Work? *International Journal of Human-Computer Studies*, 45, 185-213.

Scaife, M., Rogers, Y. Aldrich, F. & Davies, M. (1997) Designing for or Designing With? Informant Design for Interactive Learning Environments. In: *Proceedings of CHI 97, Human Factors in Computing Systems*, Atlanta, March 1997, 343-350. N.Y., Addison Wesley.

Sims, R. (1994). Seven levels of interactivity: implications for the development of multimedia education and training. In (Ed) M. Ryan, *Proceedings of the Asia Pacific Information Technology in Training and Education (APITITE)*

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.

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.

Underwood, J. (1997) Editorial: if a picture is worth a thousand words, what is the value of two? *Journal of Information Technology for Teacher Education*, 6(1), 1-6.

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