
7

External Cognition, Innovative Technologies, and Effective Learning

Mike Scaife
Yvonne Rogers¹
University of Sussex

Bruner (1972) once observed that “What a culture does to assist the development of the powers of mind of its members is, in effect, to provide amplification systems to which human beings, equipped with appropriate skills, can link themselves” (p. 53). Bruner’s (1972) thesis was to develop an account of how a society should proceed in presenting the developing child with skills or beliefs or knowledge in “a form capable of being mastered by a beginner” (p. 53). This question, although posed more than 30 years ago, is highly relevant to much contemporary Information and Communication Technology (ICT) research. Further, Bruner’s analysis is predicated on an important fact about the ways that humans operate in the world: They are highly resourceful at exploiting their environment to extend their cognitive capabilities, and they do this with a variety of strategies, tools, and representations. This, broadly speaking, is what we refer to as “external cognition.” Understanding why and how this works in the case of rapidly evolving technologies such as ICT requires, we argue, a framework (or frameworks) that allow people to see technologies in as wide a context as possible and consequently, to better understand how and why we might make use of them in a

¹Now at Indiana University

range of particular contexts: education, work, play, or leisure. In this respect, one needs to look at individual technologies both (a) as tokens of a wider type such as "tools" and (b) in terms of their specific properties or capabilities that might allow genuinely novel opportunities for learning. A result of this should be to derive some lessons for a more fruitful relation between research, design, and implementation in educational settings.

Here we are interested in developing a conceptual framework that can lead to a better understanding of the basis on which to design and use new technologies to support learning. Our focus is on the ways that technologies can allow new forms of representations—one not possible with existing media—and how these might be exploited for learning. This is not easily done simply by extrapolating from our understanding of existing media and technologies because not a great deal of generalizable information is known about the cognitive mechanisms that underpin learning from a wide variety of them, be they traditional (e.g., pictures) or novel (e.g., animations, multimedia, virtual reality; Scaife & Rogers, 1996). Thus, we proposed an emphasis on the need for an identification of the different kinds of cognitive benefits that particular representational formats and technologies may provide, what Scaife and Rogers (1996) call an analysis of "cognitive interactivity." In this chapter, we exemplify how such analysis can be used to help design learning environments by reference to research carried out by us at the INTERACT Laboratory. The examples will include a relatively mature technology, interactive multimedia (IMM), as well as two newer arrivals—virtual worlds with autonomous agents and mixed reality environments.

IMM

In terms of potential benefits for learners, IMM seems to offer much in the way of novel forms of representations (through its ability to combine graphics, animations, text, audio, video, etc.) and interactivity (by allowing the user to select, manipulate, and combine these representations). In particular, one of the major advantages offered for designing learning environments is IMM's capacity to develop novel ways of providing/viewing/interacting with multiple representations. This is a significant feature because 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, data may be shown as a table, graph, and histogram all depicting the same mathematical information in different formats. There is, here, an assumption that learning about a domain through multiple representations (different views on the same topic or concept) can engender different ideas while also helping learners constrain their interpretations. Furthermore, it is often argued that the more appro-

priately different representations are integrated by the learner, the more likely it is that “deeper” understanding will occur (Kaput, 1989).

A key research question this raises is what is the best way to coordinate different kinds of representations to support learning? One of the reasons why learners may find it difficult to integrate multiple representations is the fact that it requires additional work: They have to both interpret the individual representations and to translate between them (cf. Ainsworth, 1999; Cox, 1999; Narayanan & Hegarty, 2000; Rogers, 1999). Furthermore, the mappings between the representations may not be at all obvious to learners, making it confusing for them to try to switch between them. A key issue, therefore, when providing multiple representations for learning is to consider how to support better the translation process so that it is more obvious and explicit to the learner. This raises issues about what should be represented. For example, one strategy might be to use the representations to depict the same concept but at different levels of abstraction. This, perhaps, would lead to deeper understanding because deeper levels of description were being provided. An example of this, which we use in the system described following, is that of combining a “realistic” simulation of a domain with more abstract formalisms representing it. Used in this way, each representation can provide a different perspective but also map onto each other, guiding the learner to reflect on the relation between them. However, simply displaying representations at different levels of abstraction will not by itself enable the learner to understand the relation between them (Jones & Scaife, 2000; Rogers, 1999). What is also needed is a way of allowing the learner to actively explore the mappings.

To achieve some progress with this, we decided to apply a cognitive interactivity analysis (Scaife & Rogers, 1996). At the top level, the framework focuses on the properties of external representations in terms of their contributions to “computational off-loading”: 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 (and see Cheng, 1999; Larkin & Simon, 1987). High off-loading is where much of the effort is off-loaded onto the representation, requiring minimal effort on behalf of the learner for a given task. In our analysis, we identified several ways of achieving off-loading by manipulating the representational format (e.g., Cheng, Lowe, & Scaife, 2001; Rogers & Scaife, 1998). An apt example in the present context is that of rerepresentation—how different external representations that have the same abstract structure make problem solving easier or more difficult (see also Peterson, 1994; Zhang & Norman, 1994). These kinds of cognitive characterizations provide us with a starting point from which to begin to think about the value of using and combining different representations for different tasks. However, we need a more detailed analysis to decide, for example, on which combinations of representations to

use to aid a given stage of learning. We now outline our own research to show how such decisions may be made.

INTERACTIVE ECOLOGY LEARNING (EcoI): UNDERSTANDING DYNAMIC (ECO)SYSTEMS

The chosen area for our research was that of understanding dynamic systems. Previous work has shown that children have considerable difficulty understanding how they work (e.g., Griffiths & Grant, 1985). Here we selected ecosystems as the example domain. Younger children encounter problems in learning about this using existing materials, particularly static diagrams in which the conventions of reading the food web's interrelation using the arrow links between nodes remains mysterious. Consequently, children are unable to reason about the food web as a whole and when asked to predict the knock-on effects of extinction of a species for others in the food web are unable to do so.

To us (Rogers and Scaife, 1998), therefore, the learning task seemed to be one of (a) explicating the meaning of the food-web links and (b) supporting integration of knowledge of the parts of the food web into a meaningful whole. We felt that the learner needed to understand what the diagram is about by mapping its structure to the world it represents. IMM seemed promising here because it has the capacity for providing a means by which multiple representations can be explicitly and dynamically linked with each other ("dynamlinking"), a form of rerepresentational off-loading. This allows us to address the issue of the learner mapping between representations in ways that traditional media simply cannot. Thus, we designed EcoI, an IMM system centered on a software prototype called "PondWorld," representing a simple pond ecosystem. The basic display was an animated simulation of the creatures in the pond and another, overlaid window showing some formalism of the ecosystem. It was aimed specifically at explaining food-web diagrams for 9- to 11-year-old children who were learning about basic ecosystem concepts in school. In terms of our framework, we wanted learning to be incremental, systematically decreasing the amount of off-loading by increasing the level of abstraction and task complexity as the learner progressed. Thus, the system was designed to be experienced as a succession of modules for each of which we chose forms of interactivity suitable for its degree of difficulty and the learning task. The modules are shown in Fig. 7.1.

Module 1: Learning Factual Knowledge

The level of computational off-loading in module one is high: There is little inferencing to do, and learners have only to interact with the animation by pointing, clicking, and listening to voices. The learning process supported in this module is obtaining factual knowledge: feeding relation between a set of organisms in the community. The animation shows fish predators eat-




FIG. 7.1. Finding out about PondWorld.

ing water beetles, water beetles eating tadpoles, and tadpoles consuming weeds. The child interacts with the animation by rolling the cursor over the organisms, which produces a loudspeaker icon. Clicking on this results in a “voice” from the creature telling the learner what it is and what it eats. The voices used were designed to vary in pitch, from low (top predator) to high (weeds)—a design idea based on a suggestion by the children during an informant design session (see Scaife, Rogers, Aldrich, & Davies, 1997). For example, clicking on the largest predator results in a deep voice: “I’m a perch, I eat beetles and tadpoles.” By contrast, the primary food source says in a high-pitched, squeaky voice, “I’m a weed. I make my own food.” The child’s attention is further drawn to salient relations by having a red circle appear around feeding episodes, such as around a beetle eating a tadpole. After interacting with this module, a multiple choice quiz is presented on the screen allowing the learner to immediately test their knowledge of the feeding relations.

Module 2 “IntroWeb”: Learning to Read the Diagram

Here the level of computational off-loading is still relatively high. The learning process supported in this module consists of understanding what

the canonical notation (the arrow) used in food-web diagrams represents, that is, that the species at the head of an arrow eats the one at the tail of the arrow. In this module, the learner is presented with two adjacent representations: a food-web diagram and a concrete simulation of it. The two representations are coupled using dynalinking: The organisms in the animation are designed to behave in relation to the abstract feeding relations depicted in the food web. For example, clicking on the arrow link between the weed and tadpole in the food-web diagram results in the animation showing a token weed slowly being eaten by the tadpole. The learner has to select different feeding relations (as represented by the arrows) in the food-web formalism and observe the outcomes of their action in the concrete animation. After familiarizing themselves with the task by clicking on some of the arrows, children were asked to make predictions about what would happen in the pond before clicking on other links.

Module 3 "LinkWeb": Constructing the Diagram

Here the level of off-loading is reduced, as the learner has to consolidate and generalize knowledge from the previous module by constructing a food-web diagram in a new situation in which extra species are added. The child fills in the links by clicking on the organisms in the correct order (e.g., slime eaten by snail) on the diagram. Again, the purpose of dynalinking was to encourage the learner to make explicit links between the different levels of abstraction. Feedback is displayed in the form of colored arrows, which appear when the correct feeding relations have been linked. Children were asked to complete the whole diagram.

Module 4 "EraserWeb": Holistic Reasoning

The final module 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, that is, when one of the species is removed. The objective here was to get the learner to reason about the ecosystem as a whole (i.e., what the consequences will be for the other organisms) by reading off and interacting with the food-web diagram. An example perturbation is demonstrated initially through a narrated animation of what would happen to the ecosystem when the tadpoles are removed. The demonstration shows crosses placed sequentially on the organisms in the diagram to indicate the knock-on effects through the ecosystem as a consequence of tadpole extinction. At the same time, their concrete counterparts in the adjacent simulation are removed from the pond—the module again emphasizing the dynalinking between the two forms of representation. The child 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. To complete this, the child is required to drag and drop crosses from an adjacent palette onto the organisms they think will die in the food-web diagram, having verbally made predictions of what would happen first. Hence, the level of off-loading 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 reason correctly about which species will or will not be plausibly affected.

EVALUATING ECOI

The final version of Pondworld was tested with 9- to 10-year-old children working in pairs to see whether it was effective for supporting learning in the ways that we hypothesized. The children were allowed to interact with the software with a minimum of intervention by the adult experimenter and took from 25 to 45 min. Pairs were given a pretest and a posttest to assess learning. The chief benchmark for success was that experience with the software should result in a generalizable understanding of the abstract formalism of the food web such that the child will be able to reason better about possible changes to the ecosystem purely on the basis of a diagram. The formal testing, before and after using the system, utilized a different food-web diagram to PondWorld to factor out the possible influence of rote knowledge (or clever guesses) of species behavior. The children were asked a series of questions about the different species tokens in the web such as (a) "what does this one eat?," requiring that they read the arrows correctly; (b) "what will happen if we take this one away?," requiring that they identify the knock-on effects of removal in a simple way; and (c) "what produces its own food?," a test of understanding how a food web layout usually places consumers and producers at different levels.

Overall, 11 pairs out of 14 (79%) showed statistically significant improvement and 3 pairs (21%) showed none. This was good evidence of improvement, but more revealing findings came from the detailed behavioral observations during their experiences with PondWorld. We saw many instances of an "aha!" learning experience when the results of the animation contradicted experience. For example, on the IntroWeb module, 1 pair's dialogue was as follows when discussing what would happen if they clicked on link between species:

- A: "Nothing will happen because the weed makes its own food" (initial prediction)
- B: "Oh No! The tadpole's eaten the weed" (as a result of seeing the animation). And then a generalization to a previously misunderstood interaction with a simple three-item food chain seen in class

- A: "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."

By the time they had got to EraserWeb, most of the pairs were 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 weed being removed—indicating that they were understanding and using the diagram in the way it was supposed to be used. This contrasts sharply with their inability in the pretest. As an example, one pair discussed the implications of an organism eating more than one other organism. The system presents the scenario: The weeds have gone from the pond.

- C: 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).
- D: but that (pointing to the snail) eats weed as well.
- C: That eats weed as well, so Yeah, but that eats slime as well, so that can survive on slime.

The apparent success of Ecoi would seem to demonstrate that our insistence on and methods for mapping between cognitive analysis and design were well founded. We argued that the use of dynamically linked multiple representations, abstraction tied to simulation, would facilitate the semantic task of understanding what the diagram links were about. We further argued that progressively increasing the level of abstraction would help to integrate understanding into a more holistic domain model. Finally, we proposed that allowing the learner to actively manipulate the representation would be important. All of these claims seem to have been supported in what we observed. Indeed, the children actually performed on the posttests at a level not expected of them until several more years of schooling had been completed.

VIRTUAL WORLDS, VIRTUAL CHARACTERS

An IMM system such as Ecoi can, as we have shown, provide a rich environment in which new forms of representation can be offered to learners to give them a different view of difficult problem spaces. However the kinds of phenomenological experiences are limited when compared to more powerful three-dimensional (3D) virtual environments (VEs) that can offer opportunities for users for novel interactions with objects and navigation in 3D space. Many claims have been made about the benefits of interacting with these kinds of VEs such as easier learning, better understanding, and greater engagement (e.g., Allison, Wills, Bowman, Wineman, & Hodges,

1997; Psozka, 1995). More specifically, Wickens (1992) proposed that VEs encourage people to be more active in the way they interact with external representations through having to continuously choose their position and viewing perspective when moving through the virtual environment. In so doing, Wickens suggested that learning and retention of information can be increased. Whatever the general truth of such claims, it is the case that for our research, just as for the IMM example of the previous section, undertaking development of VEs necessitates posing specific questions about the particular properties of the media being implemented in relation to the task/goals under consideration. In this part of the chapter, we briefly describe an example of this from our own research (see Scaife & Rogers, 2001, for more details).

BUILDING A VIRTUAL THEATRE

The example project, the Virtual Puppet Theatre, had as its goal the creation of a virtual theatre for young (4–8 years old) children to support learning through playing (see Marshall, Rogers, & Scaife, 2002). Here, our aim was to provide young children with a means of extending their existing capacity for storytelling by providing them with a new set of tools that they could use to create, edit, direct, and act out plays in a virtual, imaginary setting. To do this, we created a child-friendly 3D world modeling a farm that allowed real-time interaction with and between believable life-like agents (virtual characters such as a farmer, sheep, or cow; see Fig. 7.2). The value of believable agents for storymaking has already been demonstrated, for example, by the Stanford Virtual Theatre and the Carnegie Mellon University (Mateas, 1997) Oz Project who developed the Woggle agent for micro-worlds. There is also increasing use of agents as assistants for children's learning in complex domains, for example, the Intellimedia project (Lester, Stone, & Stelling, 1999). For us, the value of using agents was first, that their behavioral characteristics could be scripted, for example, the response they exhibit when encountering another agent; and second, that they were autonomous, that is they would move and act by themselves. Thus, a simple story line could be set up and acted out by the agents on the basis of how their parameters were set. In the basic version of the virtual puppet theatre the child could, therefore, watch the story unfold as the agents encounter and interact with each other. However, such passive viewing was not particularly engaging for the children and does not provide much in the way of learning opportunities. For this reason, we built further facilities into the system that allowed the child far more interactivity.

One of our claims about the potential value of a virtual setting over and above a physical one is that it can provide an extensive range of novel supports for children to be creative. We started from a particular point of view on the




FIG. 7.2. A scene from the virtual theatre showing a low-status cow and a high-status farmer.

possible differences between play in a physical (real world) and virtual environment. During the normal course of improvisational play, the child is dealing with ephemeral actions/interactions: They occur and are gone, and the child has no time to reflect on them. Using our (Scaife & Rogers, 1996) ideas on externalization, we argued that by contrast a virtual theatre can be structured to assist development of a number of different skills. For example, autonomous agents offer the possibility of characters, inside play, which are not ephemeral (in the child's imagination). They offer the potential for the child to "read" the motivations and intentions of the characters in situations, which are still playful. Providing avatar facilities, allowing them to view the world through the eyes of individual characters, is of particular interest here. The cognitive benefits of all this, we hypothesized, was that reading of agents will be a useful means for exposing the child to situations in which such decentering skills can be polished (cf. Piaget, 1972).

The scripting for the agents was based on a dramatic scenario of them trying to achieve certain simple goals: The farmer wished to have the animals neatly in their stables/pens; some animals wished to stray and others could

act to either help or hinder the farmer. Goals were mediated through setting parameters for the states of the agents. The first was “status” (high or low), which determined how dominant or submissive the agent was in the interaction. The second, “attitude” (positive or negative) characterized the strategy of each, for example, cajoling versus herding for the farmer, fleeing versus confronting for the cow (see Fig. 7.2). These parameters changed as a function of the continuing interaction: They were not fixed. Thus, the basic model for the world was quite simple, but the combinatorial possibilities ensured that up to 10 min of varied interaction was possible between any two characters for a single run.

However, we wished to specifically promote the child’s reflection on and consequent understanding of the different roles involved in story development and enactment. This is achievable by having the virtual theatre provide a variety of editing tools to build up/change agents’ personality traits and behavior. For example, we provided a touchscreen interface that showed a number of icons portraying, for example, happy or sad expressions. Selecting an icon set the starting state for the agent in question for the next run of the scenario. Another facility allowed the child to stop the action and to record voices for the character. These clips then become attached to the particular status/attitude combination that the agent was demonstrating at the time of recording and would be played whenever the state was revisited during the scenario. Having such interactivities available allows, therefore, a continuum between the situation of just observing, as when children watch agents behave within the 3D world, and ones in which they manipulate the situation in a more direct fashion. The child could thus move between

- Spectating: The child observes the agents act in the virtual environment.
- Acting: The child joins the stage as an avatar and acts together with the agents.
- Directing: The child can make choices in terms of the agents’ character or appearance and edit elements in a story line.

EVALUATION OF THE VIRTUAL PUPPET THEATRE

We have carried out an extensive evaluation of the virtual theatre (Marshall et al., 2002). A main finding was that the children’s ability to read unambiguously what the agents are “up to” within the dramatic scenario described previously is not clear-cut. In observing mode, they do not seem to read the intentions and behaviors of the agents with much accuracy. This seems, however, to be at least partly due to their reduced interest—they have expectations derived from playing computer games that seem to determine what they perceive, for example, the farmer is often glossed as trying to

kill the cow. However, once there is a greater degree of involvement (interactivity) such as taking an avatar view or being involved in recording, they pay far more attention, and the legibility of the characters increases markedly. This, however, has a strong age effect. Younger children (4 or 5 years) disregarded the status/attitude cues in deciding what sound each character should make. For example, when the farmer was in the state low-status, positive attitude, one 5-year-old chose to record as a sound file "I'm mad because of you! [angry tone]" and as a sound for the cow, also in a low-positive state after being taken back to its pen by the farmer, "I won't stay in my kennel [defiantly]." However, even if young children showed little sensitivity to the expressive elements of the improvisational scenario, this provides evidence that they were able to grasp the basic narrative and to some extent play with it.

Older children, however, exhibited a much greater sensitivity to the scenario. For example, an 8-year-old girl, when asked to provide a sound for the cow in low-positive, recorded "Do I have to? Will I still get my supper if I stay out here? I promise I really will be really good Mr. Farmer. Please." Thus, older children are sensitive to status differences. The same girl also demonstrated an understanding of character viewpoint, deducing that the sound made by the farmer when standing alone must be a monologue: "I think he must be talking to himself because he can't be talking to the cow," and "I'm glad Daisy's gone back in her pen now. She's a good girl and she shall get her supper."

The children seemed to greatly enjoy their encounters with the virtual puppet theatre and its characters once it was sufficiently interactive. The great successes of the project seems to lie in the demonstration of the power of agent worlds to engage children in playful interactions and of the value of providing tools for externalization to develop their understanding of motives and intentions.

MIXED REALITY ENVIRONMENTS (MRES)

Both the virtual puppet theatre and Ecoi were projects that explicitly focused on the power of representing things in a novel way in a virtual form. However, a different set of challenges and possibilities are offered by more recent advances in the design of interactive technologies, namely, MREs. Drascic and Milgram (1996) offered one description of them: "Between the extremes of real life and Virtual Reality lies the spectrum of Mixed Reality, in which views of the real world are combined in some proportion with views of a virtual environment" (p. 123). A good example of this is the work at the Nottingham Mixed Reality Laboratory (2001) where a poet was immersed in a virtual environment to control a virtual representation of himself and see a video view of the audience displayed within the world looking

out from the stage. From a theoretical point of view, we consider a potential distinction for such MREs as being that between (a) a real world where spaces and artifacts are acted on by conventional physical actions and where the user's understanding is, therefore, in terms of general causal models of the world, and (b) a virtual world where a different and as yet little-understood set of causal models are applied and action may be arbitrarily coupled to the properties of the perceived world. However, the scope of MREs as sketched out previously is already insufficient. We now also have the possibility of extending the ontological profusion of worlds and objects to include in our MREs artifacts that might appear like regular physical objects but that have embedded intelligence of some sort: ubiquitous computing devices (Weiser, 1993). This raises the question of how people will deal with MREs that combine real, virtual, and ubiquitous forms. Such environments provide a challenge for understanding how we—as users—might appropriate them but also offer possibilities for learning in new ways and just as important, learning about new things. Thus, in the last of our three examples, we focus on the possibilities of MREs.

CONCEPTUALIZING MRES

The kinds of example ICT that the reader has already encountered in this chapter, IMM and VEs, have been glossed in terms of learning by using much conventional psychological terminology and concepts. However, in thinking about MREs, we need to take a different tack. Understanding what children, for example, will make of an MRE—and making design decisions about how such environments can be best configured—requires a high level of description in terms of a framework, at least initially. In the INTERACT Laboratory, we began by conceptualizing the issue in terms of “experiences”: What is the phenomenology of being in and interacting with an MRE? How can we describe this or design for it? We have begun, as a heuristic ploy, to develop a taxonomy for describing a mixed reality experience in terms of “transforms”: changes in the state of the world (Rogers, Scaife, Gabrielli, Smith, & Harris, 2002). People encounter and represent transforms between states of the world routinely in everyday life, for example, in perception (e.g., seeing an object disappear/reappear or changing one's viewpoint), in action (e.g., when the purpose of a gesture changes), and in cognition (as when we rerepresent and reinterpret the state of the world). Dealing with transforms will involve the user in some implicit or explicit theory of what causes changes of states, and this is something that we can investigate.

In an MRE, we can have real, virtual, and digitally enhanced objects or spaces, and we can design experiences that can cross between them, for example, action with a physical device, such as a wand, could result in an effect

in a virtual space such as a projected display. We have decided to use the term *digital* (D) to refer to actions/activities/effects involving virtual or digitally enhanced artifacts. For the remainder we use *physical* (P) as a cover-all term. It is important to note that in using D and P as labels, we are referring to the mechanism that potentiates the transform. Thus, a "D action" (such as using a painting program on a display screen) inevitably involves some degree of P action on the part of the user, but it is the (D) mechanism that allows this that is crucial for us.

THE CHROMARIUM: AN MRE FOR EXPERIMENTATION (2002)

To try out these conceptual distinctions we designed an MRE for young children as part of our contribution to the Equator Project (Rogers et al., 2002a). This was called the Chromarium (a space where color may be contained, observed, and experimented on). The goal here was to create a mixed reality space that enables children to experiment and play with color mixing across different media and representations. Thus, we designed a series of color-mixing tasks that could utilize either P or D actions and have their outputs in either P or D forms. This 2×2 design therefore gave us four classes of activity/transform types. For example, the P→P transform, with action and effect of the same kind, was mixing paints in pots. The D→P transform, with action and effect of different kinds, used interaction with an animation on a computer screen used to cause the differently colored blades on a small physical model windmill to rotate (mixing their colors).

Pairs of children, aged 6, took part in the Chromarium study. They were told that they would be mixing colors in fun and unusual ways and were allowed to explore the activities as they liked. To illustrate how the Chromarium was implemented, we described one example, the P→D transform. Here we used Radio Frequency (RF) technology to enable physical actions to trigger a virtual effect (RF tags are devices that can be hidden in objects, broadcasting their presence and thus recognizable to a suitable reader). Two colored blocks having a different color displayed on each face were built, each with an RF tag inside. The children could select blocks and place them on a table surface where they were detected by the tag reader concealed beneath the table. An animation (written in Macromedia Director, Version 8, 2001) mirroring the colors would appear projected onto a wall showing the effect of mixing the selected colors accompanied by a variety of sounds (Fig. 7.3).

We wanted children to both experience and reflect on their interactions, allowing some insight into their conceptualizations of these environments. The results from the children's explorations were quite complex. They en-




FIG. 7.3. Children exploring the physical-digital transforming in the chromarium.

joyed activating the various transform types and were able to collaboratively discover creative ways of using the mixed reality setups for the activity of color mixing. They also made a variety of experiments into and reflections on the novel causality of the MRE. For the $P \rightarrow D$ transform, the example described previously, the children began by placing the blocks in towers to see whether anything would happen (we had not designed for this combination). They also pressed the blocks down hard on the table surface as if trying to amplify or speed up the feedback from the animation. Some of the children tried out different exploratory behaviors to discover how the block faces were read. They tried to put the blocks against the computer screen that was behind them or against the image projected on the wall to see if “any effect was coming out.” They also tried to see if they could select a block’s face for the animation by orienting it toward the area projected on the wall.

Activities performed in the mixed conditions, $P \rightarrow D$ and $D \rightarrow P$, turned out to be the most effective for prompting children to provide causal explanations about the mechanism enabling the transforms. Their understanding of the causality involved in these transforms seemed to be affected by the fact that their actions produced effects in a different space. When we asked the children to explain where the effects (animations) of the physical manipulation of the blocks came from, they said things such as “the effect is

coming from the computer's screen over there, and it arrives here by means of electricity." Another child pointing, first at the table under which the tag reader was concealed, then to the projected image, then to the computer, said "... connect, connect, connect ... wire ... it is connected under here [table] and goes all the way up to there!" [PC behind him]. Thus, children seemed to understand that there were causal links between distant objects and their actions, but their explanations were often an idiosyncratic mix of magical thought with bits of previous knowledge. For the D→P condition, some children did not believe that they could control in the virtual space the effect obtained in the physical world. A 6-year-old girl, for example, said that the spinning of the physical windmill was caused by the wind coming in from a window (even though no open window was available to justify this). Other children, however, were much more likely to seek contingent reasons for causal relations. For example, one pair looked under the table to discover which device was causing an effect and how the different bits of kit were connected to each other.

THE HUNTING OF THE SNARK: AN MRE FOR EXPLORATION AND IMAGINATION (2002)

The Chromarium was a first attempt to design MREs for the specific purposes of examining children's causal models of the new technologies and to see how these might be adapted for science experimentation. The Hunting of the Snark project has a wider brief, being aimed at enabling young children (8–10 years old) to create and be part of novel mixed reality experiences both indoors and, eventually, outdoors (Rogers et al., 2002b). This MRE is named as a nod to the Lewis Carroll story poem in which a group of people describe their encounters with an elusive, never really found, fantasy creature, the Snark. In our MRE, groups of children are given the task of finding out about a Snark: its form, likes/dislikes, personality, and so forth. The hunt made use of a search space—to look for things to help them find out more about the creature—and a separate set of other spaces, each instantiating a different MRE, as areas to go and look for it. The searching for clues was done using a handheld computer or "snooper" that allowed the children to move around a room and discover hidden objects. The snooper was sensitive to X-Y location in the room and could show a variety of visualizations on its screen as an aid to discovery. The objects that the children discovered could be used to enter or activate a space in which the Snark might be. Here we describe one example of such a space: "the well."

Initially, the children use the snooper to find "food" for the Snark. These are plastic food facsimiles (e.g., tomato, onion, piece of chicken) with an RF tag embedded inside each. The children take the food to the well, a place where the Snark is known to sometimes be. The well is, in essence, a hori-

zontal screen allowing back projection so that visualizations may be projected on to it. At the well (Fig. 7.4) we provided a feeding device, a slot, where food could be posted into the (virtual) water that could be seen rippling at its surface. As the children approached the well, there was a sound of snoring coming from it. They could then rouse the Snark by gestures (interpreted through a gesture recognition system) and feed it by dropping food through the slot (with appropriate visualization of food dropping, sound of a splash, etc.). The result of these actions was to trigger different visualizations on the screen. For example, if the Snark was vegetarian (we could change its nature), vegetables would be received with a sight of an appreciative mouth and the sound of happy feeding noises. Feeding it meat would result in disgusted noises and an angry mouth. Once the food was all gone, the Snark would appear to swirl in the well and disappear through a tunnel in its side, cuing the next search for clues and discovery.

We used fairly abstract and changeable visualizations and sound effects to maximize the imaginative potential of the MRE. The children were told that they had to discover as much about the Snark as possible during the hunt, and this would be difficult if a highly specific and unchanging creature was provided. We gave them a “camera” to capture information, explaining that this would capture sound as well as visuals. It was basically a vibrating physical toy to which we added a “lens” and capture button. The camera worked by sending a signal to the software infrastructure when the




FIG. 7.4. The “Well” showing a happy feeding expression.

button was pressed, allowing the system to store the current state of the Snark (what the children could see/hear of it at that point). At the end of the hunt, they could “download” what had been taken with the camera and see/hear it again. This procedure instantly reminded them of their experiences with the Snark, enabling them to reflect on what it was like, what it was doing, and so forth.

We have, here, only described the well, but there are several other activities in place involving the use of a range of technologies such as wearable computing and GPS devices. The children who have tried out the prototype system have been highly engaged and were fascinated by their MRE experiences (Rogers et al., 2002b). For the well, they had no difficulty understanding what they needed to do with the physical food tokens they had collected, that is, to transform them into “electronic” food for the Snark to eat, and they found the effects of their actions, the Snark liking or disliking certain foods, to be compelling and instantly recognizable. The children would reflect on these preferences, often based on their own experiences. For example, one child said the Snark didn’t like onions was because it made it cry—as was true for the child.

In the Chromarium, we tried to conceptualize the activities in terms of transform types. For Snark, a key additional concept is that of “traversals”—experiencing things that cross the boundaries between physical and virtual spaces. Thus, the food objects in the well are physically present one moment and virtual food objects the next: They cross an ontological boundary. There is surely much potential here for the imaginative design of a new generation of learning environments.

CLOSING THOUGHTS

We began with a reminder of Bruner’s (1972) insight into the value of technologies as amplifiers for bringing learners into the culture. In this chapter, we have encountered a richness in terms of ICT that must place it firmly among the ranks of such amplifiers with its potential to improve learning and understanding of the world. Our emphasis throughout has been that it is an extremely good idea to go about realizing this potential by first putting in place a theoretical framework that allows one to map between the possibilities afforded by technology and the kinds of skills or learning or other experiences one wishes for the child. Such a mapping will necessarily be context dependent, varying with the particular technology (what it might or might not do) and the goals for which we might use it. We summarize our own experiences in this regard as well as the lessons of the exercise.

For IMM we were able to focus on the ways in which dynalinking could “bring to life” the conventions and semantics inherent in a particular class of abstractions. With PondWorld, we hoped that the interactivity provided by

the program would be effective in allowing students to go beyond the information given (Bruner & Anglin, 1973) to generate hypotheses to fill in the gaps in their understanding. However, we needed the cognitive interactivity framework to be able to design the specific modules: choosing an appropriate interactivity to gradually reduce the degree of computational off-loading. In this case, the evidence is quite clear that there was an increased ability to handle abstraction and to use the abstraction to make inferences about the represented domain.

The virtual world (agent/3D) technology offers a different set of potential properties to IMM, in particular, new kinds of objects to interact with. Again, in designing the system, we looked for a form of externalization that is consistent with our emphasis on the value of external cognition and the system's ability to do something that is difficult or impossible to do in the everyday world. For the virtual puppet theatre, we were able to use the capacity to interact with the autonomous agents as a platform for bringing out children's often implicit ideas about the motives and intentions of others and various editing facilities to allow the child to "step back" from the immediate action, promoting reflection and perhaps the development of narrative skills. This is, potentially, a highly creative process enabling children to do something relatively novel and one with a potentially substantial learning benefit as well.

The final class of technologies we looked at, the MREs, presented a less clear picture in the search for a mapping between cognitive analysis, technological capability, and learning advantages. One reason for this is that some of the basic elements of the MRE are still very much under development, for example, the various wireless technologies. Another is that one cannot clearly formulate cognitive analyses of optimal use in learning contexts until we have a better idea as to how children might appropriate the technology in the first place. Thus some of our research effort has been about discovering the ways children interact with and form causal explanations for the new experiences available in an MRE. The preliminary investigations reported here have shown that children's play and creativity as well as their attempts at scientific investigation can be well supported by digitally augmented objects that children are free to manipulate in a 3D space as representational or pragmatic devices. Of particular relevance in this context is how children's creativity appears to be enhanced when their physical activity in the everyday world leads to rich multimedia effects in a digital space.

In this chapter, we have tried to show how cognitive interactivity might be a valuable conceptual tool, certainly for designing IMM and virtual world applications. We have also discussed an initial set of taxonomic devices for MREs, that of transforms and traversals, in an effort to capture and design for a new kind of experience for the user. It is these kinds of development that are needed to drive the design of effective learning environments.

ACKNOWLEDGMENTS

The authors gratefully acknowledge support for the research cited here from (a) the ESRC Cognitive Engineering Programme, Award No. L127251033 for the Ecoi project; (b) The EU i3 ESE Programme, Project EP 29335 for the Virtual Theatre project; and (c) The EPSRC for the Equator project, Award No. GR/N15986/01. Finally, thanks to the children, schools, and parents in East Sussex, United Kingdom for their help.

REFERENCES

- Ainsworth, S. (1999). The functions of multiple representations. *Computers & Education*, 33, 131–152.
- Allison, D., Wills, B., Bowman, D., Wineman, J., & Hodges, L. (1997). The virtual reality gorilla exhibit. *IEEE Computer Graphics and Applications*, 17, 30–38.
- Bruner, J. S. (1972). *The relevance of education*. Middlesex, England: Harmondsworth.
- Bruner, J. S., & Anglin, J. M. (1973). *Beyond the information given: Studies in the psychology of knowing*. London: Allen & Unwin.
- Cheng, P. (1999). Interactive law encoding diagrams for learning and instruction. *Learning and Instruction*, 9, 309–325.
- Cheng, P., Lowe, R., & Scaife, M. (2001). Cognitive science approaches to understanding diagrammatic representations. *Artificial Intelligence Review*, 15, 79–94.
- Cox, R. (1999). Representation construction, externalised cognition and individual differences. *Learning and Instruction*, 9, 343–363.
- Drascic, D., & Milgram, P. (1996). Perceptual issues in augmented reality. In M. T. Bolas, S. S. Fisher, & J. O. Merritt (Eds.), *SPIE Volume 2653: Stereoscopic displays and virtual reality systems III*. San Jose, California, USA, 123–134.
- Equator Project. Retrieved October 10, 2001 from www.equator.ac.uk
- Griffiths, A. K., & Grant, B. A. (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.
- Jones, S., & Scaife, M. (2000). Animated diagrams: An investigation into the cognitive effects of using animation to illustrate dynamic processes. In M. Anderson & P. Cheng (Eds.), *Theory and applications of diagrams. Lecture Notes in Artificial Intelligence*, 1889, 231–244.
- Kaput, J. J. (1989). Linking representations in the symbol systems of algebra. In S. Wagner & C. Kieran (Eds.), *Research issues in the learning and teaching of algebra* (pp. 167–195). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Larkin, J. H., & Simon, H. A. (1987). Why a diagram is (sometimes) worth ten thousand words. *Cognitive Science*, 11, 65–100.
- Lester, J., Stone, B., & Stelling, G. (1999). Lifelike pedagogical agents for mixed-initiative problem solving in constructivist learning environments. *User Modelling and User-Adapted Interaction*, 9, 1–44.
- Marshall, P., Rogers, Y., & Scaife, M. (2002). The value of a virtual environment for learning about narrative. *ACM SIGGROUP Bulletin*, 23(2), 14–15.
- Mateas, M. (1997). An Oz-Centric Technical Report CMU-CD-97-156. School of Computer Science, Carnegie Mellon University, Pittsburgh, PA.

- Mixed Reality Laboratory. (2001). Retrieved October 11, 2001 from www.mrl.nott.ac.uk/projects/performances.html
- Narayanan, N. H., & Hegarty, M. (2000). Communicating dynamic behaviors: Are interactive multimedia presentations better than static mixed-mode presentations. *Proceedings of the 1st International Conference on the Theory and Applications of Diagrams*, Springer-Verlag, pp. 178–193.
- Oz Project. Retrieved September 5, 2000, from www-2.cs.cmu.edu/afs/cs.cmu.edu/project/oz/
- Peterson, D. (1994). Re-representation and emergent information in three cases of problem solving. In T. H. Dartnall (Ed.), *Artificial Intelligence and Creativity* (pp. 81–92). Dordrecht: Honer.
- Piaget, J. (1972). *The child and reality*. Paris: Editions Denoel.
- Psotka, J. (1995). Immersive tutoring systems: Virtual reality and education and training. *Instructional Science*, 23, 405–431.
- Rogers, Y. (1999). What is different about interactive graphical representations? *Learning and Instruction*, 9, 419–425.
- Rogers, Y., & Scaife, M. (1998). How can interactive multimedia facilitate learning? In J. Lee (Ed.), *Proceedings of First International Workshop on Intelligence and Multimodality in Multimedia Interfaces* [CD-ROM]. Menlo Park, CA: AAAI Press.
- Rogers, Y., Scaife, M., Gabrielli, S., Smith, H., & Harris, E. (2002a). A conceptual framework for mixed reality environments: Designing novel learning activities for young children. *Presence: Teleoperators and Virtual Environments*, 11, 677–686.
- Rogers, Y., Scaife, M., Harris, E., Phelps, T., Price, S., Smith, H., Muller, H., Randall, C., Moss, A., Taylor, I., Siznton, D., O'Malley, C., Corke, G., & Gabrielli, S. (2002b). Things aren't what they seem to be: Innovation through technology inspiration. In *Proc. Dis '2002, Designing Interactive Systems Conference*, AMC, pp. 373–379.
- 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. (2001). Informing the design of a virtual environment to support learning in children. *International Journal of Human-Computer Studies*, 55, 115–143.
- Scaife, M., Rogers, Y., Aldrich, F., & Davies, M. (1997). Designing for or designing with? Informant design for interactive learning environments. In S. Pemberton (Ed.), *Proceedings of CHI'97* (pp. 343–350). New York: ACM Press.
- Virtual Theatre Project. Retrieved September 5, 2001 from www.ksl.stanford.edu/projects/cait/
- Weiser, M. (1993). Ubiquitous computing. *Computer*, 26, 71–77.
- Wickens, C. (1992). Virtual reality and education. In *Proceedings of IEEE International Conference on Systems, Man, and Cybernetics*, 842–847.
- Zhang, J., & Norman, D. A. (1994). Representations in distributed cognitive tasks. *Cognitive Science*, 18, 87–122.