International Journal of Artificial Intelligence in Education (2000), 11, 177-196

The Fragmentation of Culture, Learning, Teaching and Technology: Implications for the Artificial Intelligence in Education Research Agenda in 2010

Gordon McCalla *Department of Computer Science, University of Saskatchewan, 57 Campus Drive, Saskatoon, Saskatchewan S7N 5A9, CANADA*[1]
e-mail: mccalla@cs.usask.ca

Abstract. My goal in this paper is to try to characterize Artificial Intelligence in Education (AIEd) research a decade hence. By then, the increasing universality of information technology will have so overloaded people with information that they will find it necessary to drastically constrain their interactions in cyberspace. The result will be a major trend to *localization*, not globalization. This localization will have two main aspects, both resulting in a fragmented social environment. The first is that people will live in their own personal *electronic villages*, and will view cyberspace locally from there, accessing only that information and contacting only those other people that are consistent with their own perspectives and goals. The second is that cyberspace will be partitioned into a massive number of *virtual communities* each with a global geographic reach but a narrow conceptual focus. People in their villages will be members of only a few such communities. Knowledge will flow relatively slowly from community to community, impacting people only when it enters their village through the communities in which they participate. This will, of course, have major impact on the nature of learning and teaching, which, in turn, will affect the AIEd research agenda. The issues the field considers to be important, the kinds of technology that it builds, even the way research is carried out, may all be transformed.

INTRODUCTION: THE UNIVERSALITY OF INFORMATION TECHNOLOGY

By 2010, information systems will be widely and deeply integrated into peoples' lives. Continued rapid growth in the capability of computer and communications technology and the software that gives it functionality will have ensured that computation will be part of most daily activities, radically affecting many occupations, changing the world of play and entertainment, stimulating major paradigm shifts in the arts and sciences, fundamentally altering how we view ourselves in relationship to the world. The trends are already obvious. The "information revolution" is the subject of numerous articles, books, and media reports. Bell Labs predicts the world will be enveloped in an "global communications skin" by the year 2025, into which everybody will be wired and which will track us and interact with us continuously (Bonisteel 1999). In all the hype, it is important to realize that the information revolution has only just begun. Ultimately, the effects will be at least as significant to human society as the renaissance, the industrial revolution, and the invention of the printing press put together.

As a field on the cutting edge of current technology, Artificial Intelligence in Education (AIEd) will continue over the next decade to be well positioned to understand the implications of (and contribute to) the information revolution. Moreover, since learning is so central to most human activities, AIEd research will increasingly tackle a wider and wider range of problems affecting every aspect of life. AIEd's unique concerns with deep computational models of learning (distinguishing it from other areas of learning technology) will also be advantageous when such deep modelling will be necessary to achieve the functionality required of learning

[1] The first draft of this paper was written while the author was a Visiting Scholar at the Basser Department of Computer Science, University of Sydney, Australia, 1998-99.

technology in 2010. As today, this should give the field the right "forcing functions" (echoing Brown, 1990) to develop new ideas that have much broader relevance than just to AIEd.

In this paper, I will consider the implications for the AIEd research agenda of the changes that will be coming over the next decade. First, I take a very general perspective, examining at a social level how the increasing universality of information technology will affect the spread of knowledge through society, and, correspondingly, the role of learning. I then take a much finer grained view, looking at the evolution of learning and teaching towards a model where learning happens incrementally and just in time. This is followed by a brief overview of the information technology that might be around in 2010 and how this will affect the kinds of systems AIEd researchers will be building. Having described the social and technological context in 2010, I then specifically focus on AIEd, discussing what will be the field's driving forces, and how the current main research issues will be transformed. The goal will not be to predict what AIEd will have learned over the next decade since in many ways this is impossible (otherwise, I'd be writing journal papers that would be ten years ahead of their time!) My goal is instead to predict how the field will have changed, what it will be investigating in 2010. The paper then concludes by reiterating the main driving forces for AIEd and by reinforcing the field's unique position to contribute broadly and deeply to the information revolution.

Before proceeding with my analysis of the brave new world of AIEd research in 2010, I will make one prediction that I am *sure* will come true: there will be at least one major technological or social development in the next decade that nobody right now is predicting, but that will have huge unanticipated implications for society. Much as the web created brand new research issues and whole new industries, this major development will have an enormous impact on the world, on technological directions, and on AIEd research issues. With this cautionary note, let me now proceed to my detailed arguments.

FRAGMENTATION OF CULTURE: THE ELECTRONIC VILLAGE

We live in a society increasingly fuelled by information. By 2010, almost all of the world's information will be accessible through information technology. This information will include processes, text, video, other media, sensory data, in short any information gathered or used by or for society. This vast sea of information will be partitioned into a horde of web sites and other local caches. In effect, the world of 2010 will be a *billion* channel universe, far exceeding the "mere" 200 channels discussed in many of today's analyses. Of course, our real lives will still go on, with matters of nutrition, affection, family, work, religion, recreation, etc., absolutely central to us, as now. However, even in these matters, information technology is increasingly involved. By 2010, it will be deeply embedded in resolving any information need, personal, social, or commercial.

Such universality of information technology will not bring a universal world view, a global culture. The exact opposite, in fact: information technology is likely to be a force for *localization*. Living in a world with ubiquitous and unavoidable access to information means that people will increasingly have to restrict access to this information, in order to maintain control of their lives. They will only allow access to information that relates to their own interests (professional and personal). They will receive input mainly from people they know or who they feel might have something worthwhile to offer. They will block most other input. In effect, each person will be at the nexus of a limited range of electronically accessible people, information, and other resources that form a sort of *electronic village*. In marked contrast to the predicted global village of the television era to be shared by everyone, this particular village in cyberspace will be definitively local. It will exist for no one else: it is a relativistic construct that exists only from the point of view of the person who "lives" there. To adapt Turkle's phrase, we will all be at "the center of our many stories" (quoted in McCorduck, 1996).

Much as in a real village, this electronic village will constrain a person's perspectives, giving them only a fragmentary and partial view of the wider world outside the village. Within his or her electronic village a person will access different information and contact different people in order to fulfill different needs. If the person wants information about the next AIEd conference, they will consult the AIEd Society website; if they want to know when their next dance class is, they will send e-mail to another member of the dance class; if they want to prepare for an exam in their software engineering course, they will get in touch with their course study group. These other people (and this other information) will be viewed by the person as playing certain roles, achieving certain functions, just as in a real village the baker supplies bread, the butcher provides meat, town meetings are announced on a notice board outside the town hall, or a neighbour acts as sounding board for ideas. While a person will likely maintain a fairly stable village environment as a basis for their perspectives on cyberspace, nevertheless, as in a real village, a person's electronic village will be subject to change and to outside influence. New people can move into the village (or at least travel through), current neighbours can move out, people can leave the village and return with new perspectives, outside information can filter in through e-mail and other means. Of course, all of this movement is virtual and relative to the person whose village it is. This person can add friends to their village, stop interacting with others, look more widely into cyberspace or talk to friends who have done so, receive e-mail or other electronic contact from afar, etc.

Thus, inevitably, new ideas will filter into a person's village. Even if something new comes to the attention of the person, however, its significance and meaning may well elude them without the help of others in the village. The person may often only be able to interpret a new idea in depth, realize its importance, appropriately contextualize it, in short turn it into *knowledge*, when somebody else in their village can help them to do so. In turn, once a person has come to understand something, they will be able to help others in their village understand it in depth. This re-inforces the centrality of localization, rather than globalization, in the information-rich world of 2010. While the *flow of information* may be governed by the speed of light, the *flow of knowledge* will be governed by the speed of human to human interaction. This will be one the keys to learning and teaching in 2010.

FRAGMENTATION OF CULTURE: VIRTUAL COMMUNITIES

Within their village a person will be a member of many *virtual communities* (in the sense of Rheingold, 1998) that extend far beyond the boundaries of their village. The other people in their village, i.e. the person's neighbours, will also variously be members of some of these same communities. However, the person won't personally know most members of these communities. Moreover, he or she will not be a member of most communities in cyberspace. Figure 1 shows the relationship of a person's own electronic village to the virtual communities intersecting and surrounding their village.

Some of the communities will be explicit, for example the community that is the AIEd Society. Some of these communities will be implicit, emerging as informal patterns of interaction on the internet and connectivity on the web. In fact, work by Kumar, Raghavan, Rajagopalan and Tomkins (1999) shows that there are already over 100,000 such "emerging communities" on the web, not yet even acknowledged but still definable as clusters of related interests. Some of these will later evolve into explicitly defined communities, others will fade away, and still others will stay in an implicit state. Membership in a virtual community will be fluid, with members coming and going with ease. While each community will be geographically distributed, it will be conceptually focussed around a few issues; for example the community of people who believe in a flat earth, the community consisting of users of a particular piece of software, the HO-gauge model railroad enthusiasts, Canadian Computer Science Department Heads, the international AIEd research community, etc.

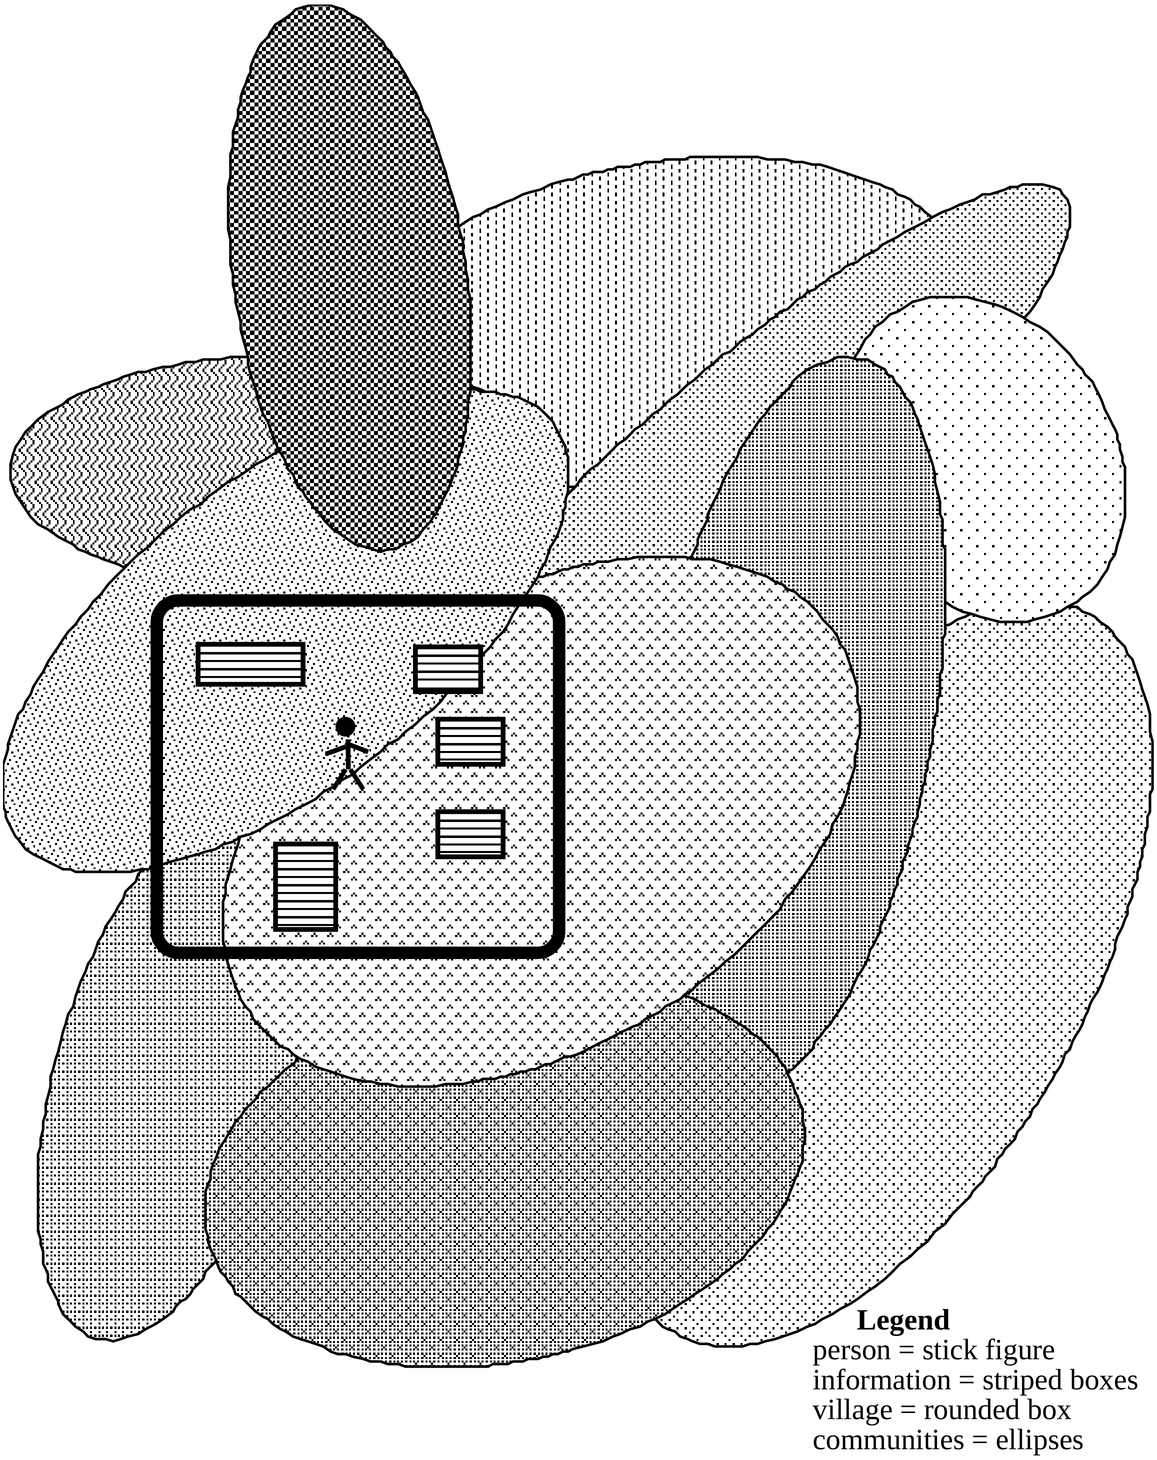

Figure 1. - A Person's Electronic Village and Surrounding Virtual Communities

The importance of learning will vary from community to community. Some of these communities will be explicitly educational, defined in terms of traditional learning goals and structures, such as the community consisting of the students taking Cmpt. 115 at the University of Saskatchewan. Some will be extensions to such traditional learning situations, such as a set of people with varying background and skills living throughout Australia who want to learn about Windows 09. Some communities may be non-educational in goals, but have a learning component, such as the Society of Professional Engineers, where it is important to learn many facets of Engineering practice through "apprenticeship" to senior members of the community. Some may not be perceived as educational at all, such as the tall ship enthusiasts, although much knowledge about tall ships and sailing more generally could be learned within the community. Thus, even though only some communities will be explicitly focussed on learning as a goal, learning will be a concern in *every* community culture. Moreover, the rapidity of change in all of these communities brought on, in part, by their heavy involvement with technology, will make learning and adaptation a prime concern, as it is today in any technology-based area.

The people in each community, even though they mostly will not know each other, will have common goals around a common interest. Each community will have its own terminology, its own shared assumptions, its own important issues, in short its own language and culture. The linguistic and cultural distinctness of each community will be the second major key to learning and teaching in 2010. In effect, virtual communities will structure the flow of knowledge through cyberspace, and thus will be critical to how learning happens. The following pattern should be common: relatively quick dissemination of knowledge within a community, relatively slow dissemination between communities. This is because people in a community speak the same language, share the same cultural assumptions, and frequently interact with one another. Once a few of them have understood something new, they can act as mentors and teachers in helping the others in the same community to understand it, a process that can happen fairly quickly because the new knowledge has already been translated appropriately for the rest of the community. However, for knowledge to cross community boundaries requires somebody or something to identify that it might be useful to another community, and then to do the translation, a relatively time consuming task. This, then, sets up two kinds of learning: (i) learning between communities and (ii) learning within a community.

Learning between communities

The fragmentation of cyberspace into communities will stimulate the need for learning to happen between communities, for knowledge to spread from one community to another. There will be straightforward fact-based knowledge to be disseminated, for example letting various scientific communities know about a new government policy on research funding developed in bureaucratic communities. But, there will also be much deeper kinds of knowledge that must be shared, for example fostering the spread of new computational modelling ideas from the knowledge representation research community to the commercial software vendor community or to the database research community. Different groups to whom the information is flowing may have very different learning needs. Thus, the commercial software vendor community would take a considerably different perspective on what is or is not important about a new knowledge representation technique than would the database research community. Whatever the nature of the knowledge that is spread from one community to another, it will have to be interpreted for all members of the other community in that community's own language in a way sensitive to that community's culture.

The technology built to support learning between communities will have several roles. One role will be to encourage the movement of relevant knowledge from community to community. To some degree this will happen naturally, via the spread of ideas through people whose electronic villages overlap more than one community and who take the time to transform information from one community so it is understandable to one or more of the other communities in their electronic village. Over time, new knowledge will percolate through much of cyberspace, although the spread will be slow and at each transformation it may get more and more distorted. There thus will often be a need to make a quicker, more direct connection between communities that are far apart in cyberspace, but which have knowledge to share with one another.

Descendants of today's recommender systems (CACM, 1997) may be able to help in this task. Current recommender systems are reactive to a request from a user, for example making recommendations for a person about movies he or she might enjoy, based on tracking that individual's past preferences. Future recommender systems may well be much more proactive, acting independently of a user's specific request, and seeking much more complex kinds of knowledge, often from communities far removed from those of which the user is a member. To do this, such systems will have to be continuously trawling cyberspace on behalf of their user for new knowledge. They may well form community models in order to be able to identify relevant places to look for new knowledge and to understand the language and cultural significance of any knowledge they find (see the learner modelling section, below). Research has already begun into some of the issues raised by multiple student models in group situations (Hoppe, 1995).

Once the knowledge to be spread has been identified, technology also has a role in helping to distribute it. Sometimes the technology itself will be able to spread the knowledge directly, for example having the recommender system post fact-based information to the user. But, for more complex knowledge, people may be needed to help the user to interpret the new knowledge. When the knowledge is found from a community of which the user is a member, this process will be quite easy since the user can usually be put in touch with another person from his or her own village who is also a member of the same community. There will be few linguistic or cultural problems in understanding the new knowledge and its relevance. If the knowledge is complex and is found in a community unknown to the user, the technology will have to find an appropriate person (or persons) from this other community to form the communication conduit for the knowledge. This person will act as a *knowledge broker*, a role more analogous to that of a diplomat than to a teacher, and the style of interaction would be *knowledge negotiation* (Moyse and Elsom-Cook, 1992), rather than a more traditional teacher-learner session. A good candidate to act as a broker might be a person who "lives" in the user's electronic village, and is also a member of the other community. Such a person would already know the user and they would share some mutual interests, even if these mutual interests were not focussed through the community in question. If such candidates cannot be found, then there would have to be deeper inferences to determine a compatible broker. Finding appropriate contact people from other communities to help in the dissemination of complex knowledge could be based on natural extensions of systems like PHelpS (Greer et al, 1998a), which finds ready, willing, and able peer helpers for workplace learning.

Once the knowledge broker has been identified, collaborative support tools will also be needed to foster the knowledge negotiation between them. These can be extensions of today's communication tools and collaborative learning environments (eg. see Dillenbourg and Self, 1992; Dillenbourg et al, 1995; O'Malley, 1995; among many others). The support will need to be deep or shallow, depending on the degree of commonality among the user and the broker, the amount of overlap in their villages, and the linguistic and cultural distinctness of communities in the user's village and the community from which the knowledge has been found. The fewer the commonalities, the more support will be needed in understanding terms and concepts, identifying important issues, accessing related knowledge, or even in finding further collaborators to act as "translators".

Learning within a community

Learning will also happen within communities. Following up the discussion in the previous section, an important kind of learning within a community will be to disseminate the knowledge brought in from outside. The user who has just learned something new may wish to communicate it to others in their own community. Unlike the "diplomats" who negotiate mutual understanding between communities, the collaboration here is more like a coach or a tutor interacting with students, although the "teaching" methods may be very different from traditional classroom techniques. Apprenticeship styles of teaching and learning (Collins et al, 1989) would frequently be appropriate, where the person(s) bringing the knowledge to the community will act as guides to other members of the community who need to know the information. Again, there can be support tools for this, some of them similar to the tools supporting collaboration between communities, and some different, recognizing the different relative status of the collaborators in a community to community communication and the collaborators within a community. In particular, tools to support cognitive apprenticeship, for example extensions of systems such as Sherlock (Lesgold et al, 1992), will be of value.

Another kind of learning within a community will be through internal debate, leading to new knowledge and insights. This will be achieved both through relatively free form intra-community discussion, and also "top down" through community "elders" sharing their insights with others in the community. Of course, such elders will be relativistically determined according to the issue(s) being debated, a person being an elder on some issues, and a member of the tribe on others. Technology such as MOOs and MUDs and other electronic forums, presumably extended to proactively engage individuals in discussion, will be crucial for fostering internal learning within the community.

A third kind of learning within a community involves initiating into the community newcomers who want to become members of a community. What it takes to learn to be a member of a community will depend on the nature of the community. Traditional approaches to teaching and learning, supported by "traditional" kinds of technology, may well work very effectively in the Cmpt. 115 community or the Windows 09 community, or other communities whose goals and approaches are explicitly educational and technological in current terms. For most communities, however, learning to be a member of a community will mean learning the community language and culture implicitly through immersion into the culture of the community.

Learning systems to support community immersion will have to reflect the alternative reality of the community, to use the language and terms of the community, to be embedded in interaction and pedagogical strategies that take into account the cultural assumptions of the community. Such issues are already being explored for natural human cultures by the situated learning community (eg. Lave, 1988; Brown et al, 1989; Clancey, 1997; Greeno, 1998; etc.). As Wasson (1996) has pointed out, there is nothing contradictory between the AIEd modelling perspective and the situated perspective, that is model-based learning systems can be used to support situated (or constructivist) styles of learning. Thus, AIEd will play an important role in designing support systems for community immersion. Supporting such immersion will be even easier in virtual communities than in natural human cultures, since a virtual community is so limited in scope. Moreover, its assumptions, its membership, and its knowledge will tend to be in a form that is directly accessible to the support technology. Immersion support technology should be able to identify specific learning needs of the newcomer, based on differences between their individual learner model and a model of the community as a whole. It should be able to identify collaborative partners in the community to help the newcomer, if possible people in the community who are also perceived by the newcomer to be in his or her own electronic village through mutual memberships in other communities. It should be able to provide help with terminology and assumptions in the new community. And so on.

FRAGMENTATION OF LEARNING AND TEACHING

Virtual communities thus fragment the flow of knowledge through society. Moreover, learning increasingly is on-line, contextualized by community culture, and integrated with other aspects of community life. In such an environment, there will be many opportunities for *just in time learning* of relatively small chunks of knowledge in the context of use. Thus, learning also becomes fragmented and "bite-sized" (Bonar, 1988). As is increasingly the case even today, manuals and textbooks will be thrown away, and learning will proceed incrementally while real world goals are being accomplished. Such learning will be better integrated with the learner's broader goals, and should be perceived to be more directly relevant to a learner. Learning and life start to merge, and the technologies to support learning start to merge with those deployed to support other activities in life. Already new learning technology fields are forming based on this trend, such as Computer Supported Collaborative Learning (eg. CSCL, 1997) which has spun off from the field of Computer Supported Collaborative Work. This trend will continue.

Although some learning will happen exclusively through the mediation of information technology, very often there will be a need for humans to help out, to take on the role of teacher or collaborator. As discussed in the last section, this will be especially true when knowledge is translated from one community culture to another, where human help will often be needed to interpret new knowledge from one community in ways understandable to a member of some other community. Such help will have to be reactive to fragmented learning styles. Teaching will thus also be fragmented. This style of just-in-time teaching has much more in common with tutoring and coaching than with traditional one-on-many teaching (i.e. with the "studio" rather than the "transmission" style in Andriessen and Sandberg's, 1999, terms). This may stimulate a re-birth in the original goals of intelligent tutoring systems (ITS) to provide one-on-one, learner-centred, reactive support for learning, but transformed into supporting not just the learner, but also the teacher (as is being explored now by Kumar, McCalla and Greer, 1999).

It should be noted that an important part of the fragmentation process is the fragmentation of roles. That is, everybody will potentially be either teacher or learner, depending on the situation. Of course, learning occurs in either role, as has been pointed out by research into "reciprocal teaching" (Palthepu et al, 1991; Nichols, 1993). It will also be important to evolve current coaching, cognitive apprenticeship, tutoring, and teaching styles to fit into the new fragmented teaching and learning context. Education research has always been a valuable source of theories and empirical results (as discussed by Winne, 1992, Shuell, 1992, and many others) for AIEd. As education research also begins to shift paradigms to more fully explore fragmented styles of teaching and learning (rather than classroom styles), then AIEd should be able to learn even more from education.

It is possible to identify two kinds of technologies that will be crucial in such a fragmented learning and teaching environment. One will be "help" technologies that stay in the background and come into play when a learning need is identified. Early versions of such technologies have already started to appear (Greer et al, 1998b). These technologies will have to determine when there is a need to intervene, and then design a pedagogical strategy to help the learner once intervention has happened. The need to intervene can be externally generated by people or technology, as when somebody brings knowledge to a community, or a recommender or filtering system suggests something may be of interest. Such a need can also be internally generated by procedures that monitor individuals as they use information technology to determine whether they need help, much as today the "intelligent" helper supports a user of Microsoft Office software. Such monitoring will be carried out using the most advanced diagnosis techniques of the day (eg. much as Bayes Nets are used today in the current Microsoft Office helper, see Horvitz, 1997). This diagnosis will be able to be more sophisticated than now, however, due to the availability of a plethora of data about the learner accumulated because so many of their activities have been carried out on-line. The learner should thus be free to carry out normal activities largely unaware of the monitoring. There will be no need to hold the learner to a fine-grained model as in the early model-tracing tutors (Anderson and Reiser, 1985). Once a learning need has been recognized, the pedagogical strategies designed by the learning technology can be similarly well informed, in particular taking advantage of knowledge of the culture and language of the community in which the learning is happening. There will be further discussion of these things in the section below on AIEd research issues in 2010.

The other kind of AIEd technology that will support learning in such a fragmented world will be simulation and discovery environments. Currently, such environments are viewed as special purpose learning tools, applicable in very narrow situations. It is thus hard to see how they can be integrated into the more general learning landscape. However, in a fragmented world, such environments have a natural home. They can be imported into particular communities, to serve particular learning needs, once these are identified. Discovery environments will be especially useful when there is an identified need for particular members of a community to master difficult skills, requiring them to synthesize knowledge (rather than to just learn facts). Systems to support learners as they practise skills in such environments will be the inheritors of much from current learning technology research, perhaps somewhat surprisingly the work in both discovery environments and standard tutoring systems. The discovery environments, suitably enhanced by the latest in virtual reality and multimedia, will provide a realistic microworld, as in many current AIEd research projects (eg. Sherlock, Lesgold et al, 1992; BioWorld, Lajoie, 1993; Smithtown, Shute and Glaser, 1990). The tutoring systems, enhanced by the latest developments in AIEd, will be able to provide deep guidance, in the mould of Elsom-Cook's (1990) guided discovery tutoring approach. Even if they impose restrictions on the learner, this will not be seen to be inappropriate because the learner will recognize the imposition as being only temporary, to serve the fragmented learning needs of the moment.

SOFTWARE WITHOUT BOUNDARIES

In addition to the impact on AIEd caused by the fragmentation of culture, learning, and teaching, the changes in software technology itself will affect what it means to design and build an AIEd system. Even today, it is becoming ever more difficult to actually define what is and isn't part of a particular piece of software, what its boundaries are, what it does and doesn't do. There are many factors suggesting that this trend will only accelerate over the next decade. By 2010 the concept of "a piece of software" may simply be undefined.

The rapid introduction of agent technologies will be one factor hastening this blurring of software identity. Particular software systems will be fragmented into many quasi-independent agents, and will dynamically access other agents as needed. These agents may be specifically designed for an application system or may be totally external to the application system, as software reuse becomes a reality. Advances in distributed computation and communications technology will make irrelevant the notion of a single "host" computer for these agents. Standardized high level communication protocols will make agent to agent interaction both easy and well defined. These distributed agents will interact in a way that much more resembles an ecosystem than an organized software application, with unpredictable emergent behaviours the result of complex interagent activities. Models of such computational agent societies will be drawn from the social sciences, in particular economics and sociology (eg. possibly something like actor-network theory, summarized in Ryder, 1999). But, as in these social science theories, it may be impossible to predict with precision individual effects of given actions within an application. Such work is already well underway in the agents research community (eg. Castelfranchi and Conte, 1992; Mataric, 1992) and is now being extended into AIEd contexts (eg. Vassileva et al, 1999).

There will be many other factors complicating the notion of software identity. The boundary between procedures and data will become even more blurred than it is today, with a computational entity able to take on either role as context suggests. Data will be procedurally interpreted to give it "meaning". Different procedural interpretations will give the data a different meaning, often determined only dynamically by the context of use. The concept of active data will be prevalent, with a particular information structure being able to respond actively to various queries posed to it, as in current object-oriented and agent approaches. Much of this active data will not be designed at all by humans, but will be created automatically either through direct "sensing" of the world by technological devices or through various inferences made from existing active data.

Another complicating factor in identifying a particular software entity is that it will be possible to consider any application system at many levels of detail. Coarse grained system descriptions will be able to capture an application at the social level of its interactions with the outside world, and finer grained descriptions will capture details of the system's computational properties. An application will also be viewable from many different perspectives, according to different end uses, giving it a different significance depending on which viewpoint is taken. Thus, a system has no fixed semantics or pragmatics. The best we will be able to say is that an application takes on meaning *relative* to end use, perhaps as defined by the cultural assumptions of the communities in which it is being used.

Finally, most applications will be embedded in a social context that is vastly wider than the system itself. There will many humans using the system, adding, removing, or changing data and procedures; scaffolding the system when it fails; giving orders; making mistakes; all of this possibly at the same time. The same system will be used by humans living in their own unique electronic villages, with correspondingly unique perspectives on what they are doing. These variously situated humans will be in constant interaction with the system, and the resulting feedback loops will have indeterminate effects on the system, the humans, and the culture of various communities of which the humans are members. The people who have designed and developed the system, will also have to be constantly interacting with the system in order to maintain it. It may be possible to define relativistic models of a system by taking into account some aspects of its design (as is proposed in Kay's (1999) ums interaction model) and the potential communities of its use, possibly also drawing on ideas from situated cognition (Barwise and Perry, 1983). But, the fact will remain that at a detailed level such a socially embedded system will be unpredictable both in principle and in fact. The archaic notion of a computer only doing what it is programmed to do will be finally put to rest in this complex new world of unpredictable effects, where there the notions of computer, program, and user have been fragmented into meaninglessness.

The implication of this kind of technological environment will be the impossibility of precisely defining what a particular AIEd system is and isn't. A typical AIEd application in 2010 will be drawing on resources from outside itself, both computational and human. Among other things, it will be accessing web resources (and these will be more sophisticated than today), repositories storing individual user actions, external data and knowledge bases, interaction and communication software designed for generic use, and (frequently) human input (eg. as is even now being explored in the I-Help project, Greer et al, 1998b). Moreover, the AIEd system will be designed out of lower level components that themselves may draw data and procedures from outside, perhaps even unknown at the conceptual level of the AIEd system. This outside information may be continuously and actively updated entirely independently of what the AIEd system is doing, meaning that there is no possibility of predicting in advance exactly what the system will do. It will be fundamentally non-deterministic.

Although an AIEd application system itself will be fragmented and non-deterministic, it will still be meaningful to talk about "an AIEd system", not so much in terms of its structure but in terms of its particular pedagogical philosophy. At the most general level, the goal of an AIEd application system will still be defined in 2010 much as it is today: to help the learner. A particular AIEd application system, however, will have its own specific pedagogical approach, providing its own unique coherence to the world. A system's pedagogical goals will be implemented computationally, and success in achieving these goals will be measured by comparison with effects perceived by the system in the learner(s) and in the communities in which the learning is occurring. This is pretty much how an AIEd application system would be defined today, apart from, perhaps, the emphasis on social context implied by the inclusion of "communities" in the definition.

AIED RESEARCH ISSUES IN 2010

So, in the world of 2010, people will live in their own cloistered electronic village, culture will be partitioned into virtual communities, learning and teaching will be correspondingly localized and contextualized by community culture, and information technologies will be fragmented and definable only relative to end use. It is surprising, therefore, that the basic definition of the AIEd field, as with an AIEd system, will be pretty much the same in 2010 as it is now. AIEd will still be the field devoted to exploring model-based learning technologies that are adaptive to learners, "machines who care" in Self's (1997) terms. However, the fragmentation of the technological and social environment discussed above will certainly transform the specific research issues that are important to the field.

AIEd will be an empirical science in 2010, even more so than today. The heart of the field will be building real systems for real use. To be sure there will be formal models being investigated, and there will be plenty of tool building, but these will largely spin off from discoveries made through building end use applications. "Precision in a vacuum" will be a vice! So will "technique without relevance"! Readily available will be sophisticated software engineering tools (such as interface development packages, shells for creating agent-based systems, and powerful database and knowledge base software), a plethora of "off the shelf" ready-to-use system components, and easy access to much external information on the web and elsewhere. This will make it possible to fairly quickly engineer a sophisticated prototype system for human end use.

Appropriate virtual communities will provide sets of ready made test subjects for experimenting with these systems. Learning support technology will have to be situated in the specific culture of each such community of test subjects to succeed, but if it succeeds in these communities it should be possible to adapt it to other communities. Measuring success and failure will be difficult, given the huge number of outside influences on any given community, and the idiosyncratic nature of each person's electronic village. Certainly, traditional controlled laboratory studies, with well defined control and treatment groups, will not be easy to carry out. However, AIEd should be able to borrow from discoveries in the social sciences as they, too, begin to grapple with experimentation and measurement in a culturally fragmented, but information technology rich world.

An alternative to testing with real human subjects could be to use simulated students. Despite investigations into building such students (eg. see VanLehn, Ohlsson, and Nason, 1994), it would seem that designing realistic enough students to test out the complex AIEd systems of the future will itself be an AIEd complete problem. Another possibility is the creation of standard application test beds, as has been done with the corpora used in computational linguistics and information filtering research today. Recent attempts to develop such standard applications for testing agent systems (eg. a personal travel assistant, FIPA, 1997), if successful, might show the way for AIEd applications. The problem with using such a standard test bed as a benchmark, however, is the natural tendency of a field to begin to design around the issues of the test bed, rather than those of the much broader real world, and to tailor systems to optimize performance according to metrics appropriate to the particular test bed.

In ten years time, deep computational modelling will be back in favour both inside and outside of AIEd. The cycle repeats itself! Necessary and sufficient conditions will be there to ensure this; necessary because modelling will be essential to achieve the functionality required of information systems in 2010, and sufficient because in this information-intense world it should be possible to form rich models. The modelling will be "bite-sized" reflecting the fragmentation of culture, teaching, learning, and technology. The impossibility of having complete or consistent information about the world will be obvious, and modelling techniques will reflect this reality. The models will have to be robust in the face of the unpredictability of interactions with the complex information and social world in which they are embedded. The need for realistic response times will mean that the ability to reason under resource constraints will be an essential aspect of any model. The constancy of change will be paramount, leading to lazy, context-dependent, procedural interpretation of models.

In sum, rather than having the formal logical properties of soundness, consistency, and completeness in a single declarative formalism, AIEd models will be fragmentary, robust, resource constrained, procedural, and contextualized. Fortunately, the AIEd, artificial intelligence, and computer science research communities are now headed in directions that suggest modelling techniques manifesting to a degree some or all of these properties may be available within a decade (such as anytime algorithms, genetic algorithms, Bayes nets, agent based systems, etc.). I would now like to go through various issues of traditional importance to AIEd to see what people interested in those issues might be doing in 2010.

Knowledge representation becomes knowledge management

Knowledge representation was at the heart of the AIEd enterprise for the first two decades of the field, but has fallen out of favour more recently. By 2010, it will be back in a big way, but considerably transformed. Domain knowledge representation will still be an issue, but so will representation of the learner (to capture aspects of their electronic village) and representation of social context (to capture aspects of various community cultures). There will be no distinct boundaries as to what is being represented: the domain will elide into other related domains, one community will overlap with many other communities, the learner model will be part of a much larger individual user model and will often merge into group models based on community cultural assumptions. An AIEd system will have access to much of the world's information, as needed for its activities. In fact, the specific knowledge representation goal of an AIEd system will be to filter the vast amount of knowledge, information, and process that will be readily available.

Achieving this goal will change the main knowledge representation focus from representation of new knowledge to the *management* of existing knowledge. Activities such as knowledge finding, knowledge integration, sense making, and change tracking will predominate. *Knowledge finding* will be aimed at the discovery of knowledge that is relevant to a person (and through him or her possibly an entire community), using advanced and possibly pedagogically specific search engines (such as the recommender systems discussed above). *Knowledge integration* will involve trying to manage the disparate sources of information flowing into an AIEd system, reconciling local inconsistencies, assessing credibility, etc. *Sense making* will be the related task of interpreting this knowledge according to a person's current learning goals and/or current community culture. Techniques from data mining and machine learning may well prove fruitful sources of insight in this task. Finally, *change tracking* will be concerned with keeping an AIEd system abreast of the constantly changing information universe. There will be tradeoffs between importing knowledge for use by the AIEd system and storing it "locally" for later use, on the one hand, versus seeking and interpreting knowledge "on the fly", on the other. Either way, the particular learning goals at the time of use as contextualized by the learner's own electronic village and the community culture(s) in which the learning is situated will determine the meaning of the knowledge, that is which aspects are considered to be important and which not, which interpretation to give to conflicting knowledge, etc.

AIEd will not have a monopoly on these issues. Any intelligent system in 2010 will be concerned with most of them. But, AIEd will have the particular focus of managing knowledge from the point of view of learning and teaching (much as GUIDON, Clancey, 1987, re-conceptualized an existing set of expert system production rules for pedagogical purposes). This will mean a special emphasis on learner (and teacher) modelling, and on interpreting knowledge in a way compatible with human cognition. Some of the knowledge representation issues important to AIEd will be (i) dealing with "how to", not just "what", so learners can master analysis and synthesis of knowledge, as well as learn facts; (ii) developing machine learning techniques with cognitive fidelity so the models can track human styles of knowing and learning; (iii) deriving causal and episodic interpretations of events since so much human knowledge is organized in this way; (iv) understanding motivation since in many learning situations learners are not self-motivated; (v) coming up with models of forgetting as well as remembering. A start has been made on many of these issues, for example DelSoldato and DuBoulay's (1995) work on motivation and the recent workshop on emotions at the 1999 User Modeling conference (DeRosis, 1999). Hopefully, the social sciences will be able to provide insight and theories that can be readily adapted, much as ACT* (Anderson, 1983) and SOAR (Laird et al, 1987) were subsequently incorporated into AIEd applications such as the LISP Tutor (Corbett and Anderson, 1992) and ET-SOAR (Ward, 1991). In turn, AIEd may be able to act as a crucible for bringing issues of human and social cognition together with the social and computer sciences in order to forge new and interesting insights that may, in turn, reflect back to communities much wider than AIEd.

Learner modelling: from too little to too much bandwidth

Since the ability to adapt to users will be a prime feature of any intelligent information system, it is likely that general user modelling will be a very important issue by 2010. By then the almost continuous contact humans will have with information and communications technology will allow for very fine grained tracking of their activities. It will be possible for each person to have a permanent individual data repository in which is recorded their every electronic action, as well as other things such as pointers to data structures and applications they were using when undertaking these actions, records of system internal states at the time, the electronic community in which the action took place, etc. With this vastly enhanced bandwidth, user modelling will be transformed into making sense out of too much knowledge, rather than trying to make do with too little.

User modelling will divide into two main activities. Activity (i) involves understanding the basic world view of the user, essentially tracking their electronic village (i.e. their community memberships, their perceived relationships to village friends and neighbours, information about the user, basic knowledge and beliefs of the user, etc.). Activity (ii) involves keeping track of information that is important to the current end application. Of course, these activities are not wholly independent, and will draw upon one another for information. The first activity is evolutionary, keeping track of slow changes in a person's village over time. This activity will be carried out periodically by a system attached permanently to the user, perhaps the user's own personal agent (as in the I-Help system of Vassileva, et al, 1999). This personal agent would be responsible for tracking changes in the user's village, and would also act proactively on behalf of the user to find information and people of interest to the user. The second activity is extremely dynamic, and will vary considerably with each particular end application, as a person adopts a different persona for each application, perhaps even for each task within an application, much as they adopt various roles in today's electronic games and on-line discussion forums (see McCorduck, 1996). Thus, the meaning of a user's actions will have to be deciphered at the time the user is undertaking their activities (as proposed in Kay, 1999), not in advance. When the end application involves learning on the part of the user, activity (ii) will be called *learner modelling*. Since this activity will be associated with the end application itself, learner models created by an end application will exist only so long as the end application is active. There will thus be *many* learner models over a span of time as a learner moves from task to task.

The learner modelling system associated with the end application will be able to draw on information from three main sources. The first will be raw information, such as the learner's recent patterns of activity, stored in the learner's data repository. Such information may have to be massaged, synthesized, and contextualized to the needs of the current end application. The second source will be the personal agent for the learner, which will be able to provide information about things like community memberships and personal relationships, resulting from activity (i). The third source of information will be information gathered by the end application itself, such as the learner's current task and goals, whether the learning is within or between communities (and in particular which community provides the current learning context), whether other humans are involved in the learning activity or just this one, etc. The learner modelling system will be particularly interested in putting all of this information together so as to be able to diagnose learner knowledge states (i.e. cognitive diagnosis), understand learner motivations and attitudes, figure out the social context of the learning, and determine the impact of various (fine-grained, bite-sized) pedagogical strategies on the speed and effectiveness of learning.

Another important point to make about learner modelling in 2010 is the fact that collections of learners will be as important as single learners in many learning situations. As outlined in earlier sections of this paper, collaborative learning, cognitive apprenticeship, peer help, and other styles of learning involving more than one human at a time will often be central to learning. Thus, learner modelling will also have to take into account how each individual interacts with the group (as in Hoppe, 1995).

When learners are explicitly learning from or about a community of which they are not a member, their personal agent may try to consult a *community model* of that community in order to help bridge gaps between the learner's perceptions and those of the new community. Such community models are likely to emphasize cultural and social issues of importance to the community, rather than "content" issues. They are also likely to be less dynamic than individual models, evolving more slowly with changes in community culture. Finally, it should be noted that factionalization among people within a community may in fact mean that there is often a hierarchy of community models within a single community, with some globally shared values and some others shared only by particular factions.

Where do such community models come from? The learner's personal agent could create models for communities that intersect with the learner's village, based on finding patterns in the learner's activities within those communities. However, it will *not* have the information needed to create a community model for a community that is not in the learner's village. Perhaps, the personal agent could trade the community models it creates for models created by other agents for other communities, although this would suggest a uniformity in community modelling protocols that is contrary to the growing fragmentation of technology. If useful enough, however, software developers have shown their willingness to subsume total independence into globally agreed upon protocols (as in HTML, XML, and now the proposed IML). In any event, it is clear that by 2010, learner modelling begins to consider a continuum of issues ranging from the importance of individualization through to the behaviour of entire communities.

There will be a clear need for there to be some sort of "intelligent garbage collector" attached to a user's/learner's data repository. Like any garbage collector, its job will be to keep the size of the repository under control by discarding most low-level information when it is no longer needed. The intelligence will come in deciding when something is no longer needed. Further, the intelligent garbage collector will employ machine learning and data mining techniques to find patterns of user behaviour in the raw data in the repository (thus creating a sort of episodic memory, as in McCalla, 1983), and may well have a role in incorporating activity (ii) user/learner models into the data repository once they are discarded after an application is finished.

The vastly enhanced bandwidth available to the user modelling systems in 2010 will give these systems a fighting chance to really understand aspects of the user that are important to learning. Unfortunately, gross violations of privacy will also be possible. In the next decade, society at large will have to design a legal framework to protect a person's privacy while at the same time allowing systems meant to help the user to take some advantage of this increased bandwidth. Such a legal framework will be hard to design for numerous reasons, most importantly, perhaps, the difficulty in knowing how (and which) data about a person will be interpreted by the many different agents that will exist in the fragmented technological and social environment of 2010. In fact, the social issues raised by learner modelling and other aspects of the brave new technological world I am sketching in this paper are deserving of a much broader and deeper analysis in their own separate paper. This will not be that paper!

Instructional planning: adapting the environment for learning

As discussed in Vassileva and Wasson (1996), instructional planning need not be restricted to supporting rigid "instructivist" styles of learning, but can be deployed in environments where the learner is much freer to forge their own directions. A decade from now, instructional planning will be seen not in the role of "producing course plans", but as responsible for creating and maintaining an effective environment for learning, adapted to the individual needs of the learner and the learning situation. In line with the fragmented nature of learning, teaching, and society, such planning will tend to be highly reactive to current conditions and tuned to specific community culture. There may be global pedagogical principles from which the instructional planning processes can draw, but the actual planning will be reactive and local. Rather than referring to this as planning, it might be more appropriate to use the term *environment adaptation*. Various aspects of the environment will be adapted to enhance learning, as in Akhras and Self (1996). Particular explanations may be generated for the learner, visualizations of concepts may be made available to clarify concepts, relevant features derived from the learner model may be displayed for the learner to reflect upon, other peer learners or teachers may be brought into the loop and given specific tasks to work on (as in Luckin and DuBoulay, 1999), appropriate media for interaction may be chosen, and so on. The pedagogical approach will be learner centered, with any interventions more akin to tutoring or coaching than to teaching.

Content and delivery issues will still exist, but will be much more blurred than in the course planning days. It might be more accurate to scrap this distinction as misleading, and replace it with a distinction between *deep environment adaptation* (DEA) to the learner and *surface environment adaptation* (SEA) to the learner. Roughly speaking, deep environment adaptation will involve choosing goals for the learning situation and a general strategy for achieving these goals, and then guiding the learning interactions towards achieving these goals.

Surface environment adaptation will involve adapting the "surface features" of the learning environment to most effectively stimulate the achievement of the deep goals. DEA goals will be broadly construed beyond just learning domain facts, to include learning procedural, social, and cultural knowledge. There will be various kinds of learning according to the different kinds of knowledge being learned, as in Wasson's (1990) PEPE system which distinguished fact-based learning from analysis and synthesis (roughly adapting Bloom's, 1956, taxonomy). Deep goals will drive the environment adaptation, with SEA issues being seen as a means to achieving these deeper goals. Motivation will be seen as largely derived from a learner successfully gaining control over the deep issues, rather than from the particular flashy interface that the learner may be using.

In choosing deeper learning goals, the much fuller knowledge of the learner and the learning context available in 2010 will allow environment adaptation to factor in many more issues than just the prerequisite structure of content. In particular the social environment will often be critical, most prominently, of course, the virtual community or communities in which the learning is taking place. Deep environment adaptation will be able to take into account, through individual and group learner modelling, important aspects of community culture in forming particular learning goals. In many styles of learning, collaboration will be important, and deep environment adaptation will have to factor in the existence of other players besides the learner. This suggests that issues of human interaction, being explored in human-computer interaction, dialogue modelling, sociolinguistics, and even AIEd (see Baker, 1994, for example), may augment "traditional" instructional planning techniques. Investigations are already underway into interaction planning drawing from many sources and tailored to a richer information world than just natural language text (eg. DeCarolis, 1999). Thus, deep environment adaptation will incorporate generalized versions of today's dialogue planning (and plan recognition) techniques in order that the system can appropriately interact with the learner and the other humans in the learning environment. Full natural language understanding will still be beyond the "state of the art". However, coarse-grained tracking and support for the interaction will almost certainly be possible, through natural extensions to techniques already being explored. An illustrative, but by no means exhaustive, list of such techniques include using flexible interaction scripts adhered to by all participants in an interaction (as in AutoTutor, Wiemar-Hastings et al, 1998), performing robust keyword in context interpretation of natural language (as in the L2Tutor of Price et al, 1999), using latent semantic analysis as a basis for interpretation of learner's natural language expressions (Wiemar-Hastings et al, 1999), factoring in pedagogical issues in analyzing natural language dialogue (Katz et al, 1999), and detecting and using web-based argumentation patterns to understand learners' on-line interaction (Yu and Chee, 1999).

Surface environment adaptation will involve designing the visual and textual environment in which the learning goals will be achieved, drawing on the wide range of multimedia and virtual reality tools that will be available. The deep learning goals, the social context (as defined by the community or communities in which the learning is taking place), idiosyncracies of the learner and other participants in the learning situation (provided through learner modelling), may all affect the kind of surface environment created. Particular communication packages (eg. chat environments, e-mail tools, etc.), specific visualization aids (like diagrams or graphs), animations, or even discovery environments may be appropriate surface environments, depending on the learner(s) and their learning situation. These may change as the learning situation evolves and changes.

Once all the environment adaptation processes are in action, the environment the learner is using will change around him or her, often in subtle ways not even noticed by the learner. The learner should not feel that he or she is under the oppressive control of the AIEd system. The main impact should be that the learning happens more effectively and hopefully more quickly than it would have otherwise.

Control is from afar

AIEd systems historically have had a central controller through which all decisions flow. Ten years from now, however, learning technology will be so integrated with technology for doing other things, that while in some cases there will be a clear AIEd controller (eg. the environment adaptation system), very often the ultimate arbiter will not even be part of the AIEd system. The learning support tools will be called in at the behest of other systems being used by the learner and/or at the explicit request of the learner or other people who discern a learning need. These other processes and people will maintain final responsibility for the success of the learning activities, and could therefore be seen to be the "controller". (Of course, in the final analysis it is always the learner who really has the final say about this!)

For example, a worker may generate an immediate learning need to find out how to overcome a problem with a software package they are using to do some accounting at their place of work. Assistance may be provided to help the worker overcome the problem, ranging from pointers to appropriate web pages, an on-line peer help session with somebody from the accounting community in their workplace, through to an entire course on understanding the accounting package. The worker (or the package) that invoked the assistance will be responsible for determining the effectiveness of the learning, during or after the learning scenario. Another example would be a person wanting to spread new ideas to other members of a particular community requesting the support of a communication environment supplemented with learning support tools. The person spreading the ideas would be responsible for monitoring the success of the support tools, backing them up if they fail, requesting new tools if necessary, and in the end determining that the ideas have been successfully assimilated. Thus, the main control issues for AIEd will be how to integrate the AIEd system with other information technology being used by the learner, and how to interpret the success or failure of the learning episode.

Shells are breaking up

The notion of an AIEd shell that has been popular throughout the last 15 years (Suthers, 1996; Murray and Blessing, 1997, 1999; Murray, 1999) is increasingly being subsumed into research into communication protocols and ontologies that coordinate the fragmented activities of many AIEd components and learners (eg. Ritter and Koedinger, 1996; Breuker, 1997; Muhlenbrock, et al, 1998; Wasson, 1998; Vassileva, 1998). By 2010 general protocols will have been developed that guarantee that the humans and various software fragments that have come together to form a learning environment to achieve a particular set of learning goals can compatibly work together to achieve these goals. Research into enhancing these general protocols will be part of the distributed computation research community, not AIEd. The AIEd research effort currently spent on shells will thus be largely transformed a decade hence into producing reusable components that can be used widely by all sorts of AIEd systems. This will be totally compatible with the fragmented social and technological environment of the day. Processes to perform various kinds of learner model interpretation, knowledge summarization, model updating, and pedagogical interaction, etc., will all be "shrink wrapped" (as agents, perhaps) and available to be plugged in to other systems. This will allow AIEd application-oriented research to build on other work in a very direct way, through actual component sharing. It will also allow AIEd technology to be much more widely used than in just AIEd applications. For example, an AIEd agent could perform certain kinds of knowledge summarization in a database application. Of course, AIEd application systems will be able to import agents from other applications, as well, for example an inductive learning algorithm from machine learning. Thus, the boundaries between learning technology and other technology further blur.

CONCLUSION

AIEd research, like many other fields, is cyclic, with issues and techniques being developed, explored, exploited, and discarded, only to be resurrected a few years later in a different form as forces in the outside world make them relevant again. Ideally, the cycle is actually a spiral towards increasing sophistication in the issues and techniques. Similarly, the perceived relevance of a field waxes and wanes, as its issues seem to be relevant or irrelevant in the outside world. By 2010, the changes in culture, learning, teaching, and technology discussed in this paper, will once again make the goals of AIEd highly relevant to the world. The universality of information technology will mean rapid change in just about every area of human endeavour, with a corresponding need for technology-based learning to allow humans to keep up. Learning will be deeply embedded in the electronically enhanced cultures of the day, and will merge with work and life.

The AIEd research agenda will also be highly relevant. AIEd brings so many important issues together into a single place, and then so starkly illuminates them, that AIEd researchers cannot avoid working on the hard "real" problems facing technology that is to be used in the real world (although such problems can be cleverly "end run" from time to time!) By 2010, key AIEd issues will be seen to be critical far beyond the AIEd research community. Such issues include dealing with inconsistent and incomplete knowledge, individualizing systems to the needs of particular users, fostering collaborative interaction, planning under resource constraints, tracking change in users and systems, dealing with plain old hard-to-please everyday "naive" users (rather than "helpful" or "expert" users), making systems easy to use, developing and managing distributed and fragmented software systems, etc. Even today these issues are important in AIEd. They will only become more so over the next ten years.

The AIEd field is thus ideally situated to influence and contribute not only to learning technology, but also to other fields in computer science and even disciplines such as education, psychology, anthropology, and sociology. These disciplines, in turn, should be able to feed back ideas, techniques, and information to AIEd, especially as they increasingly subscribe to aspects of the computational paradigm. Knowledge will thus flow from research community to research community getting interpreted and augmented as it flows (a knowledge transmission paradigm that echoes the knowledge flow patterns predicted in this paper!) By 2010 the synergistic interactions that result should lead to sophisticated and usable systems, as well as to new insights into the nature of learning, teaching, culture, and technology.

Acknowledgements

I would like to thank Julita Vassileva and Jim Greer for their helpful comments on this paper, and Judy Kay and Sattiraju Prabhakar for their many stimulating discussions during my administrative leave in Sydney, Australia, during 1998-99. The three reviewers of the first draft of this paper had a major impact on the final version here; I would like to thank them for their constructive criticisms. I would also like to thank my graduate and summer students over the years who have provided me with many, many interesting insights on AIEd and computer science. Of course, the opinions expressed here are purely my own responsibility, and are not to be blamed on any of these people! Finally, I would like to acknowledge the financial support for my research provided by the Natural Sciences and Engineering Research Council of Canada and the Canadian TeleLearning Network of Centres of Excellence.

References

Akhras, F.N. and Self, J.A. (1996). A process-sensitive learning environment architecture. *Proceedings of the Third Intelligent Tutoring Systems Conference*, Montreal.

Anderson, J. R. (1983). *The Architecture of Cognition*. Cambridge, Massachusetts: Harvard University Press.

Anderson, J.R. and Reiser, B.J. (1985). The LISP tutor. *Byte*, *10*, 4, 159-175.

Andriessen, J. and Sandberg, J. (1999). Where is education heading and how about AI? *International Journal of Artificial Intelligence in Education*, *10*, 130-150,

Baker, M. (1994). A model for negotiation in teaching-learning dialogues. *Journal of Artificial Intelligence in Education 5* (2), 199-254.

Barwise, J. and Perry, J. (1983). *Situations and Attitudes.*

Bloom, B.S. (1956). *Taxonomy of Educational Objectives, Handbook I: Cognitive Domain*. New York: David McKay.

Bonar, J.G. (1988). *The Bite-Sized Architecture*, Technical Report, Learning Research and Development Center, University of Pittsburgh, Pittsburgh, Pennsylvania.

Bonisteel, R. (1999). "Bell Labs predicts a 'global communications skin' by 2025", *Newsbytes*, accessible from http://www.canet3.net, November.

Breuker, J. (1997). Presentation on ontologies for AIEd systems. Panel discussion, *Eighth International Conference on Artificial Intelligence in Education*, Kobe, Japan.

Brown, J.S. (1990). Toward a new epistemology for learning. In C. Frasson and G. Gauthier (Eds.), *Intelligent Tutoring Systems: At the Crossroads of Artificial Intelligence and Education* (pp. 266-282), New Jersey: Ablex.

Brown, J.S., Collins, A. and Duguid, P. (1989). Situated cognition and the culture of learning. *Educational Researcher*, *18*(1), 32-42.

CACM (1997). Special Issue on Recommender Systems, *Communications of the ACM*, *40*, 3, March.

Castelfranchi, C. and Conte, R. (1992). Emergent functionality among intelligent systems: Cooperation within and without minds. *AI & Society*, *6*, 78-93.

Clancey, W.J. (1987). *Knowledge-Based Tutoring: The GUIDON Program.* Cambridge, Massachusetts: The MIT Press.

Clancey, W.J. (1997). The conceptual nature of knowledge, situations and activity. In P.J. Feltovich, K.M. Ford and R.R. Hoffman (Eds.), *Expertise in Context* (pp. 247-291). Hillsdale, NJ: Lawrence Erlbaum.

Collins, A., Brown, J.S. and Newman, S.E. (1989). Cognitive apprenticeship: teaching students the craft of reading,writing, and mathematics. In L. B. Resnick (Ed.), *Knowing, Learning and Instruction: Essays in Honor of Robert Glaser* (pp. 453-494). Hillsdale, NJ: Lawrence Erlbaum.

Corbett, A.T. and Anderson, J.R. (1992). LISP intelligent tutoring system: research in skill acquisition. In J. H. Larkin and R. W. Chabay (Eds.), *Computer-Assisted Instruction and Intelligent Tutoring Systems: Shared Goals and Complementary Approaches* (pp. 73-109). Hillsdale, NJ: Lawrence Erlbaum.

CSCL (1997). *Proceedings of the Second International Conference on Computer Support for Collaborative Learning*, Toronto, December.

DeCarolis, B. (1999). User modeling in mixed-initiative hypertexts, *Doctoral Consortium, Seventh International Conference on User Modeling*, Banff.

Del Soldato, T. and Du Boulay, B. (1995). Formalisation and implementation of motivational tactics in tutoring systems, *Journal of Artificial Intelligence in Education, 6*, 337-378.

DeRosis, F. (1999). *Proceedings of the Workshop on Attitude, Personality, and Emotions in User-Adapted Interaction*, Seventh International Conference on User Modeling, Banff.

Dillenbourg, P. and Self, J.A. (1992). A computational approach to socially distributed cognition. *European Journal of Psychology of Education, 3* (4), 353-372.

Dillenbourg, P., Baker, M, Blaye, A, O'Malley, C. (1995). The evolution of research on collaborative learning. In P. Reimann and H. Spada (Eds.), *Learning in Humans and Machines*. London: Elsevier.

Elsom-Cook, M. (1990). Guided discovery tutoring. In M. Elsom-Cook (Ed.), *Guided Discovery Tutoring: A Framework for ICAI Research*. London: Paul Chapman.

FIPA (1997). *Foundation for Intelligent Physical Agents*, Draft Standards Specification, Part 4, Personal Travel Assistance.

Greeno, J. (1998). The situativity of knowing, learning, and research. *American Psychologist, 53*(1), 5-26.

Greer, J., McCalla, G., Collins, J., Kumar, V., Meagher, P. and Vassileva, J. (1998a). Supporting peer help and collaboration in distributed workplace environments. *International Journal of Artificial Intelligence in Education, 9*, 159-177.

Greer, J.E., McCalla, G.I., Cooke, J.E., Collins,J.A., Kumar, V.S., Bishop, A.S. and Vassileva, J.I. (1998b). Integrating cognitive tools for peer help: the intelligent intraNet peer help-desk project. In S. Lajoie (Ed.), *Computers as Cognitive Tools: The Next Generation*, Hillsdale, NJ: Lawrence Erlbaum.

Hoppe, H.-U. (1995). The use of multiple student modelling to parameterise group learning. *Proceedings of the Seventh International Conference on Artificial Intelligence in Education*, Washington, DC, 234-241.

Horvitz, E. (1997). Agents with beliefs: reflections on bayesian methods for user modeling. *Proceedings of the Sixth International Conference on User Modeling*, Sardinia, Italy, 441-443.

Katz, S., Aronis, J., and Creitz, C. (1999). Modeling pedagogical interactions with machine learning. *Proceedings of the Ninth International Conference on Artificial Intelligence in Education*, LeMans, France, 543-550.

Kay,J. (1999). *A Scrutable User Modelling Shell for User-Adapted Interaction*. Ph.D. Thesis, Basser Department of Computer Science, University of Sydney, Sydney, Australia.

Kumar, V., Raghavan, P., Rajagopalan, S., Tomkins, A. (1999), Trawling the web for emerging cyber-communities. IBM Almaden Research Center, accessible at http://www8.org/w8-papers/4a-search-mining/trawling/trawling.html.

Kumar, V., McCalla, G. and Greer, J. (1999). Helping the peer helper. *Proceedings of the Ninth International Conference on Artificial Intelligence in Education*, LeMans, France, 325-332.

Laird, J. E., Newell, A. and Rosenblum, P. (1987) SOAR: An architecture for general intelligence. *Artificial Intelligence 33*(1), 1-64.

Lajoie, S.P. (1993). Computer environments as cognitive tools for enhancing learning. In S.P. Lajoie and S.J. Derry (Eds), *Computers as Cognitive Tools* (pp. 261-288). Hillsdale, NJ: Lawrence Erlbaum.

Lave, J. (1988). *Cognition in Practice*. Cambridge: Cambridge University Press.

Lesgold A., Lajoie S., Bunzo M., and Eggan G. (1992). SHERLOCK: a coached practice environment for an electronics troubleshooting job. In J. H. Larkin and R. W. Chabay (Eds.), *Computer-Assisted Instruction and Intelligent Tutoring Systems*. Hillsdale, NJ: Lawrence Erlbaum.

Luckin, R. and Du Boulay, B. (1999). Scaffolding learner extension in the zone of proximal development, *International Journal of Artificial Intelligence in Education, 10*.

Mataric, M. (1992). Designing emergent behaviors: from local interactions to collective intelligence. In *Simulation of Adaptive Behavior 2*. MIT Press. Cambridge.

McCalla, G.I. (1983). An approach to the organization of knowledge and its use in the natural language recall task, *International Journal of Computers and Mathematics, 9*(1), pp. 201-214.

McCorduck, P. (1996). Sex, lies and avatars, a profile of Sherry Turkle. *Wired*, April 1996.

Moyse, R. and Elsom-Cook (1992). M. *Knowledge Negotiation*, London: Academic Press.

Mühlenbrock, M., Tewissen, F., Hoppe, H.-U. (1998). A framework system for intelligent support in open distributed learning environments. *International Journal of Artificial Intelligence in Education, 9*, 256-274.

Murray, T. (1999). Authoring intelligent tutoring systems: an analysis of the state of the art. *International Journal of Artificial Intelligence in Education, 10*.

Murray, T. and Blessing, S. (1997, 1999). Special issues on authoring systems for intelligent tutoring systems, *International Journal on Artificial Intelligence in Education, 8 (Part I), 10 (Part II).*

Nichols, D. (1993). Intelligent student systems: learning by teaching. *Proceedings of the Seventh International Conference on Artificial Intelligence in Education*, Washington, DC, 576.

O'Malley, C. (Ed.) (1995). *Computer-Supported Collaborative Learning*. New York: Springer-Verlag.

Palthepu, S., Greer, J.E. and McCalla, G.I. (1991). Learning by teaching. *Proceedings of the First International Conference on the Learning Sciences*, Northwestern University, Illinois, 357-363.

Price, C., McCalla, G., and Bunt, A. (1999). L2tutor: a mixed-initiative dialogue system for improving fluency. *Computer Assisted Language Learning Journal 12*(2), 83-112.

Rheingold, H. (1998). *The Virtual Community: Homesteading on the Electronic Frontier*. accessible at http://www.rheingold.com/vc/book/.

Ritter, S. and Koedinger, K. R. (1996). An architecture for plug-in tutor agents. *International Journal of Artificial Intelligence in Education 7*(3/4), 315-347.

Ryder, M. (1999). *Brief synopses of actor-network theory*. accessible from http://www.cudenver.edu/~mryder/itc/act_net_dff.html.

Self, J. (1997). Presentation on the future of AIEd. Panel discussion, *Eighth International Conference on Artificial Intelligence in Education*, Kobe, Japan.

Shuell, T.J. (1992). Designing instructional computing systems for meaningful learning. In M. Jones and P. Winne (Eds.) *Adaptive Learning Environments*. NATO ASI Series F, Vol 85 (pp. 19-54). New York: Springer-Verlag.

Shute, V. and Glaser, R. (1990). A large-scale evaluation of an intelligent discovery world: Smithtown. *Interactive Learning Environments Journal 1*, 51-77.

Suthers, D. (1996). Architectures and methods for designing cost-effective and reusable ITSs. accessible from http://advlearn.lrdc.pitt.edu/its-arch/.

VanLehn, K., Ohlsson, S., and Nason, R. (1994). Application of simulated students: an exploration, *International Journal of Artificial Intelligence in Education 5*(2), 135-175.

Vassileva, J. (1998). Goal-based autonomous social agents supporting adaptation and teaching in a distributed environment. *Proceedings of the Third International Conference on Intelligent Tutoring Systems*, San Antonio, Texas, 564-573.

Vassileva, J. and Wasson, B. (1996). Instructional planning approaches: from tutoring towards free learning. *Proceedings of European Conference on Artificial Intelligence in Education*, Lisbon, Portugal, 1-8, accessible from http://julita.usask.ca/homepage/Eaied.htm.

Vassileva, J., Greer, J., McCalla, G., Deters, R., Zapata, D., Mudgal, C. and Grant, S. (1999). A multi-agent approach to the design of peer-help environments, Proceedings of the Ninth International Conference on Artificial Intelligence in Education, LeMans, France, 38-45.

Ward, B. (1991). *ET-SOAR: An Electrostatics Tutor*. Ph.D. Thesis, Carnegie-Mellon University, Pittsburgh, Pennsylvania.

Wasson (Brecht), B. (1990). *Determining the Focus of Instruction: Content Planning for Intelligent Tutoring Systems*. Ph.D. Thesis, Department of Computational Science, University of Saskatchewan, Canada.

Wasson, B. (1996). Instructional planning and contemporary theories of learning: is this a self-contradiction? *Proceedings of the European Conference on Artificial Intelligence in Education*, Lisbon, Portugal, 23-30, accessible from http://www.ifi.uib.no/staff/barbara/papers/Euroaied96.html.

Wasson, B. (1998). Identifying coordination agents for collaborative telelearning. *International Journal of Artificial Intelligence in Education 9*, 275-299.

Wiemar-Hastings, P., Graesser, A., Harter, D. and the TRG (1998). The foundations and architecture of AutoTutor. *Proceedings of the Fourth International Conference on Intelligent Tutoring Systems*, San Antonio, Texas, 334-343.

Wiemar-Hastings, P., Wiemar-Hastings, K., Graesser, A. (1999). Improving an intelligent tutor's comprehension of students using latent semantic analysis". *Proceedings of the Ninth International Conference on Artificial Intelligence in Education*, LeMans, France, 535-542.

Winne, P.H. (1992). State-of-the-art instructional computing systems that afford instruction and bootstrap research. In M. Jones & P. Winne (Eds.), *Adaptive Learning Environments*. NATO ASI Series F, Vol 85 (pp. 349-380). New York: Springer-Verlag.

Yu, R. and Chee, Y.S. (1999). An intelligent agent in web-based argumentation. *Proceedings of the Ninth International Conference on Artificial Intelligence in Education*, LeMans, France, 551-557.