

Socially Aware Pedagogical Agents: Steps to Building a Life-Long Learning Companion



ARIES

Gord McCalla, Jim Greer

ARIES Laboratory

Department of Computer Science

University of Saskatchewan

Saskatoon, Saskatchewan CANADA

Outline


Icing



Conceptual



Contextual



Outline

- **Contextual layer**

- what research problems need to be solved to create a life-long learning companion?
- what research threads are we pursuing in the ARIES lab that relate to solving these problems?
- where can we make progress?
 - constraints provided by educational domain
 - synergy with surrounding social environment

Outline

- **Conceptual layer**

- research perspectives emerging out of ARIES
 - iHelp and its distributed environment of people, personal agents, and information
 - fragmentation, electronic villages, and communities
 - ecological approach
 - user modelling in such fragmented worlds
 - active modelling: context as purpose, person, community, and resources available
 - life long modelling: e-portfolios and user models
 - communities of learners and health of the community
 - privacy and safety

Outline

- **Icing layer**

- incremental steps to the lifelong learning companion:
 - a roadmap to success using iHelp as an illustrative example
 - weak AI vs strong AI
 - elbows, exponential change
- what are the real goals of learning
 - the role of socialization and enculturation in learning
 - how much learning is about “content”, how much about learning to be part of a community?
 - how far do we want to push the learning companion: are there limits to what it can do? are there limits to what it should do?

Contextual Layer

- Cooking up a learning companion
 - what “ready mixed ingredients” are available?
 - content - the world’s knowledge is just a Google away!
 - models of learners, teachers, content, pedagogy
 - agents to animate these models
 - are there natural advantages of the Educational Domain?
 - curricula
 - learning goals can be known / constrained
- Do we need to “solve AI” to build the lifelong learning companion?
- How much can we leverage human learning companions?



Conceptual Layer

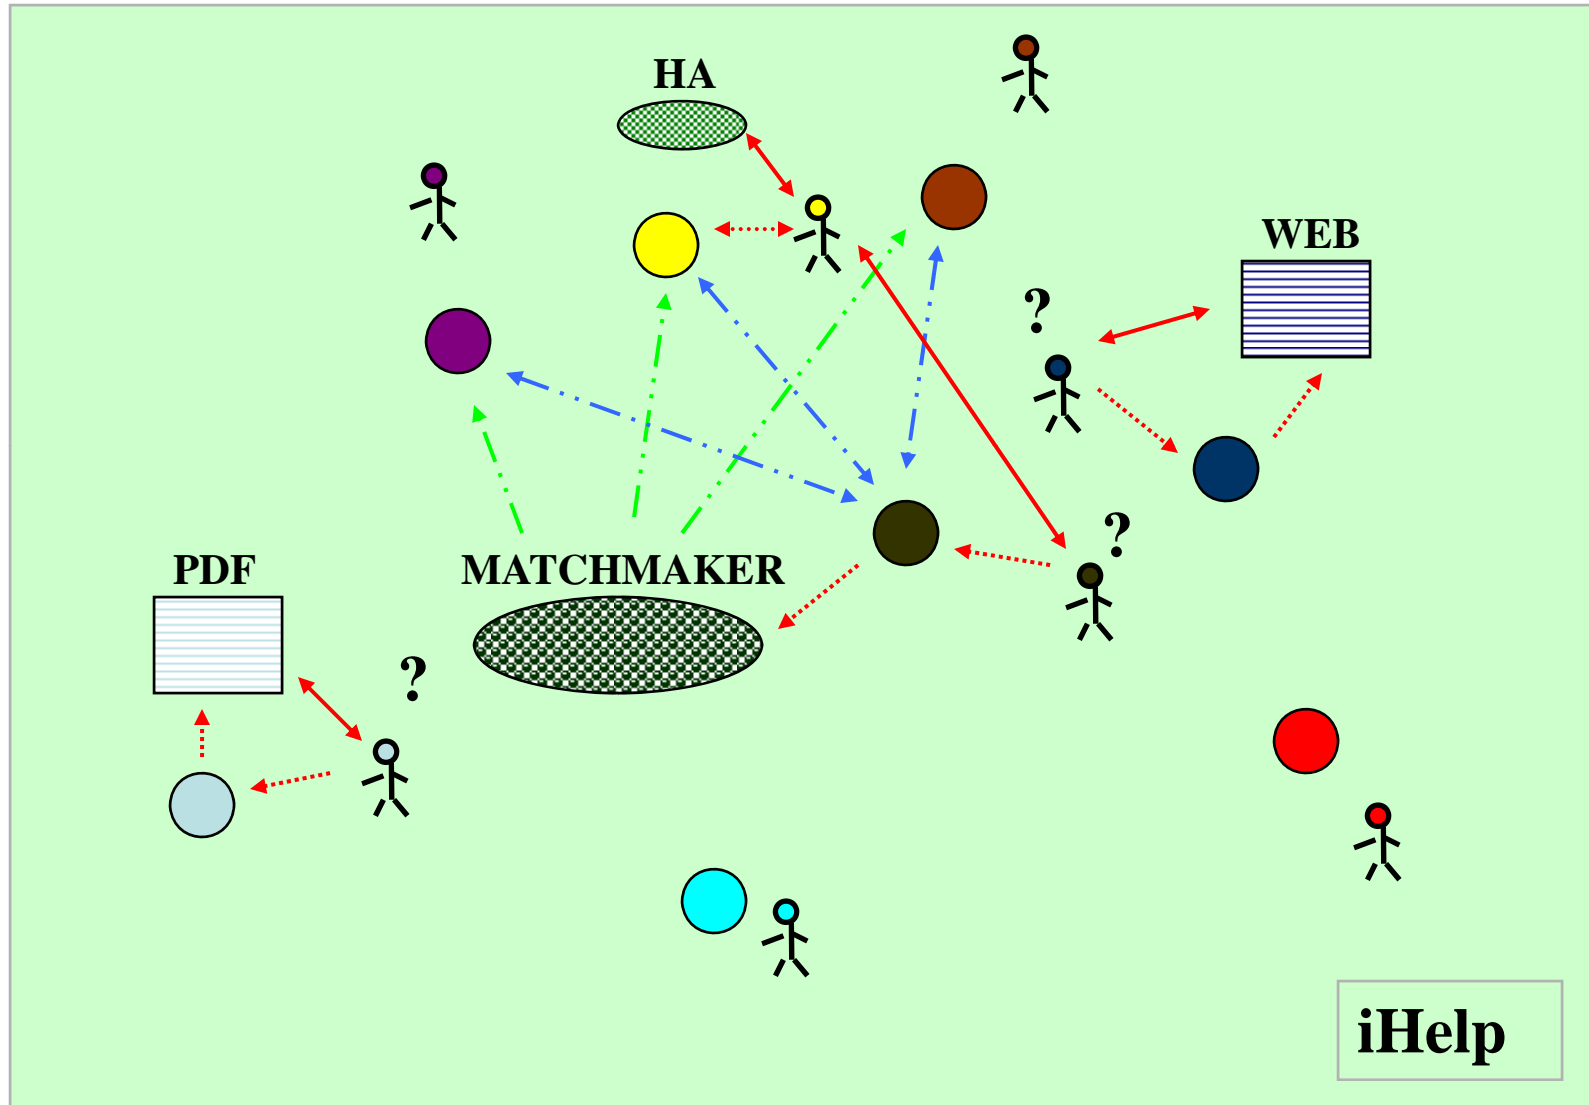


iHelp

Vision and Infrastructure

In the Beginning There was iHelp

- iHelp is a peer help system in educational and workplace environments developed in the ARIES lab in 1997 (see UM 1997 paper)
- Aimed at providing peer help and other help while learners actually solve problems, in school or the workplace: **just-in-time, contextualized**
- Underlying agent architecture in original system
 - **personal agent** for every learner, finding help and helpers to support learner - very much like a weak (very!) version of Nell's Primer
- Various versions tested in university courses and workplace, but agent architecture abandoned (for now!)
 - **thousands** of users over the last 10 years



What iHelp has Wrought

- The original iHelp has inspired a dramatic transformation in our view of what it means to support learning
 - culture, learning, teaching, and technology are **fragmented**, and learning technology must be fundamentally designed around this
 - the learning technology doesn't have to do everything: the surrounding **social community** can support the learner
 - learner modelling is **active**, "model" is a verb: the key context elements for learner modelling are the **purpose**, the **social community** (i.e. the other users), and the **resource constraints**
 - **clickstream** level data can be tracked and used by the learning technology
 - as time goes on the learning environment should naturally adapt and evolve to meet new needs: the **ecological approach**
 - **affect** is as important as content in motivating learners and teachers
 - **privacy** is important

Fragmentation of Culture, Learning, Teaching, and Technology

(McCalla, AIED Journal paper, 2000)

Fragmentation of Culture

In the billion person environment of cyberspace, people need to put up barriers to stay sane


- localized perspective on cyberspace:
 the electronic village
- each person's village is unique: relativistic, not global
- village shares many of the characteristics of a real village: neighbours, professionals, friends, community organizations, (information) markets
- person only knows something when it comes into their village

Fragmentation of Culture

Each village overlaps a wider world

- a person is also part of many **virtual communities** extending beyond their village boundaries: explicit and implicit
 - each focussed on its own issues
 - each with its own language and culture
 - overlapping each other
 - each person a member of many such communities
- learning is part of most such communities
- compare to cultures chosen in Diamond Age society

the electronic village




virtual communities

Fragmentation of Learning and Teaching

Information flows at the speed of light, knowledge at the speed of human understanding

- learning **between** communities
 - identifying the knowledge to be spread
 - supporting its spread: finding collaborators to foster understanding between communities (**diplomats**, **negotiation**)
- learning **within** a community
 - top down from community leaders and those bringing in outside knowledge (**teachers**, **apprenticeship**)
 - **collaboratively** through internal debate
 - **immersing** new village members in community culture



**learning between
communities**

Fragmentation of Learning and Teaching

Cultural fragmentation means learning is partitioned and fragmented

- “on-line”, learning can happen as needed in small chunks, in the context of on-going activities: **just-in-time** learning, fragmented knowledge
- human teachers are often needed, to help **integrate** knowledge with culture, to help **translate** knowledge into terms appropriate to learners in other communities
- each person can be teacher or learner, depending on the situation: fragmentation of **roles**

Fragmentation of Technology: Software Without Boundaries

The boundaries of a software system are increasingly indefinite

- software is fragmented into many **quasi-independent entities** (agents)
- many of these software entities come from **outside** a particular application “package” (“cloud computing”)
- behaviour of such software systems is **emergent**, like an ecosystem, fundamentally unpredictable
- distinction between **procedures and data, hardware and software** blurs
- software exists simultaneously at **many levels of detail**
- software is embedded in a complex **social environment**

Fragmentation of Technology: Software Without Boundaries

Software takes on a particular coherence only relative to end use

- as defined by the **purpose** of the technology
- as defined by the **tasks** and **goals** of the people using it
- as defined by **other people** involved
- as defined by the **communities** in which it is used
- as defined by the **resources** available

Implications for the Design of Systems to Support E-Learning Communities

- The importance of the individual
 - tools to support **personalization**
 - **user modelling**: user portfolio
 - **motivation**: the affective dimension
- The importance of communities
 - tools to support **translation** of community culture and language to other communities
 - **knowledge negotiation**, knowledge brokers
 - **community modelling**
- Centrality of context
 - focus on **pragmatics** more so than syntax, semantics
 - main context elements: **purposes, goals, tasks, users, resources**
 - nothing is independent of context: **active modelling**

Implications for the Design of Systems to Support E-Learning Communities

- Natural “forcing functions” of information and communications technology
 - fundamentally **localizing** not globalizing: fragmentation
 - broadening individual perspectives
 - **agents** as a unifying metaphor
- The constancy of change
 - tracking user behaviour: **high bandwidth of interaction**
 - possibility of **mining** user behaviour: the ecological approach
 - various purposes: recommending people or information, evaluating systems, intelligent garbage collection, ...
- Granularity
 - supporting **grain shifts**: between individual and group, between levels of knowledge

Ecological Approach

(from McCalla, JIME paper, 2004)

The Ecological Approach: Overview

- An approach to capturing information about **users' interactions** with technology and using it for various purposes
- Being explored in the context of managing **learning object repositories**
- Based on keeping **fine-grained user models** of each learner and attaching these to the learning object the learner is interacting with
- Active user modelling: **data mining**, **clustering**, and **collaborative filtering** used to make sense of the user modelling information in context: for a particular pedagogical purpose and particular learner(s)
- **Agents** representing learners and learning objects carry out the purposes

One Pedagogical Purpose

- One pedagogical purpose: recommend a learning object for a learner
 - assume: each learner's personal agent keeps a model of the person with various characteristics: personal, affective, learning style, etc.: **a learner model** (can build up over time, like an e-portfolio)
 - assume: every time a learner has interacted with a learning object a copy of the model his or her agent has kept is attached to the learning object along with a record of the interaction between the learner and the learning object: **a learner model instance**
 - then: an appropriate learning object for a learner might be one that other learners (with similar characteristics to the learner who have interacted in similar ways with similar learning objects as has the learner) have found to be useful: **collaborative filtering**
- This was Tiffany Tang's research project- the original inspiration for the ecological approach

CHARACTERISTICS






personal
affective
learning/cognitive style
previous learning objects
current goal(s)

EPISODIC

trace of learner's interactions
learner's view of content
learner's evaluation of object
outcomes

**Learner
Model
Instance**

An Example



Other Pedagogical Purposes

- There are many other pedagogical goals that could employ the ecological approach
 - to find a sequence of learning objects of relevance to a learner: *instructional planning*
 - to find out which learning objects are useful, not useful, or no longer useful: *intelligent garbage collection*
 - to find peers with appropriate characteristics: *help finding*
 - to find groups of learners with appropriate shared attributes: *building learning communities*
 - to find out what happened to a learner or learners interacting with a learning object repository: *empirical evaluation*
 - to attach relevant pedagogical information to learning objects based on actual end use experience: *metatagging*
 -

The Appeal of the Ecological Approach

- Learning objects are **activated**: they are not passive, but take on responsibilities for their use in support of learning
- Learners are “**in the loop**”: personal agents allow learners to be part of the educational environment
- Focus is on **end use**: essentially learning objects are tagged by models of the learners who use them, **not** by context-independent content tags from a pre-defined ontology of standard terms
- Approach is **ecological**: as end use experience accumulates, there can be an ever more refined understanding of what works for whom

The Appeal of the Ecological Approach

- Decision making is **contextual**: information is *actively* interpreted in context and as needed for more appropriate reactions; re-use through context capture *not* de-contextualization
- Approach is **extensible** and **adaptable**: the agent-based approach allows new learning objects and learners to be added, old ones to be deleted
- Approach is **modular**: agent approach localizes decision making and improves robustness
- Approach supports **diversity**: learners, applications, and learning objects can be integrated into one system, unified by the agent metaphor
- Potential downsides: **privacy** and **tractability**!!!!!!

Is the Ecological Approach Tractable?

- Computational issues
 - how much can be done actively
 - space-time trade-offs
 - can purposes and learner models constrain the mining, clustering, and filtering algorithms
 - can purposes cover a domain and be re-used in other domains
 - can learner models be standardized and shared
- Social issues
 - what kinds of pedagogy can be supported
 - can social capital be enhanced (Daniel, Schwier, McCalla, AIED Conference, 2005)
 - advantages of e-learning application
 - environment can be constrained
 - learner can be constrained
 - feedback from learner is natural and serves a pedagogical purpose

User Modelling in Ecological Systems

The Active Approach to User Modelling

- User modelling is a process carried out by agents
 - each agent keeps its own **models** of users/agents to help it to achieve its goals
 - these models are continuously re-computed as the agents interact with each other and with users: **active modelling**
 - each agent keeps track of its current goal and accesses only models of other agents/people implicated in helping it to achieve this goal
 - in achieving the goal, new modelling information can be computed “just-in-time” as needed and as resources allow
 - in achieving the goal, interactions with other agents and with users can also be tracked
 - the models are thus continuously updated with fine-grained information, but this information is contextualized by the goals being undertaken, the agents and people involved, and the resources available
 - an agent’s overall model of another agent or person is thus a compendium of many such purpose-based active computations
 - see Vassileva, McCalla, Greer; UMUIAI journal paper, 2003

Lifelong Learner Modelling

- **e-portfolios** and user modelling (Guo and Greer)
 - e-portfolios have different uses
 - collect, reflect, evaluate as per rubric
 - gather, organize, present to others
 - lifelong repository of artifacts, claims and evidence
 - learner models from various learning support systems can be elements of e-portfolios
 - e-portfolio is natural source of information for initializing a system's learner model
 - metadata standards do not give sufficient guidance for automating initialization
 - learner model can be built up more easily from e-portfolio (by teachers, domain experts and even by the learners themselves)

Understanding Communities of Learners

Community of Learners

- Learning technology doesn't need to stand alone: the surrounding community of learners can support the learner and the learning technology
 - need to motivate learners to participate
 - extrinsic rewards: eg. the iHelp economy (rewards in iHelp credit units - worth less even than the dollar!!)
 - intrinsic rewards: status, reputation
 - result is enhanced social capital (Daniel)
 - user modelling need only be “good enough” to improve the opportunities for learning
 - finding peer helpers in iHelp that are ready, willing, able
 - feeding back reputation to learners
- Main goal of learning may in fact be socialization into a community rather than mastery of content

Motivating Participation in Community

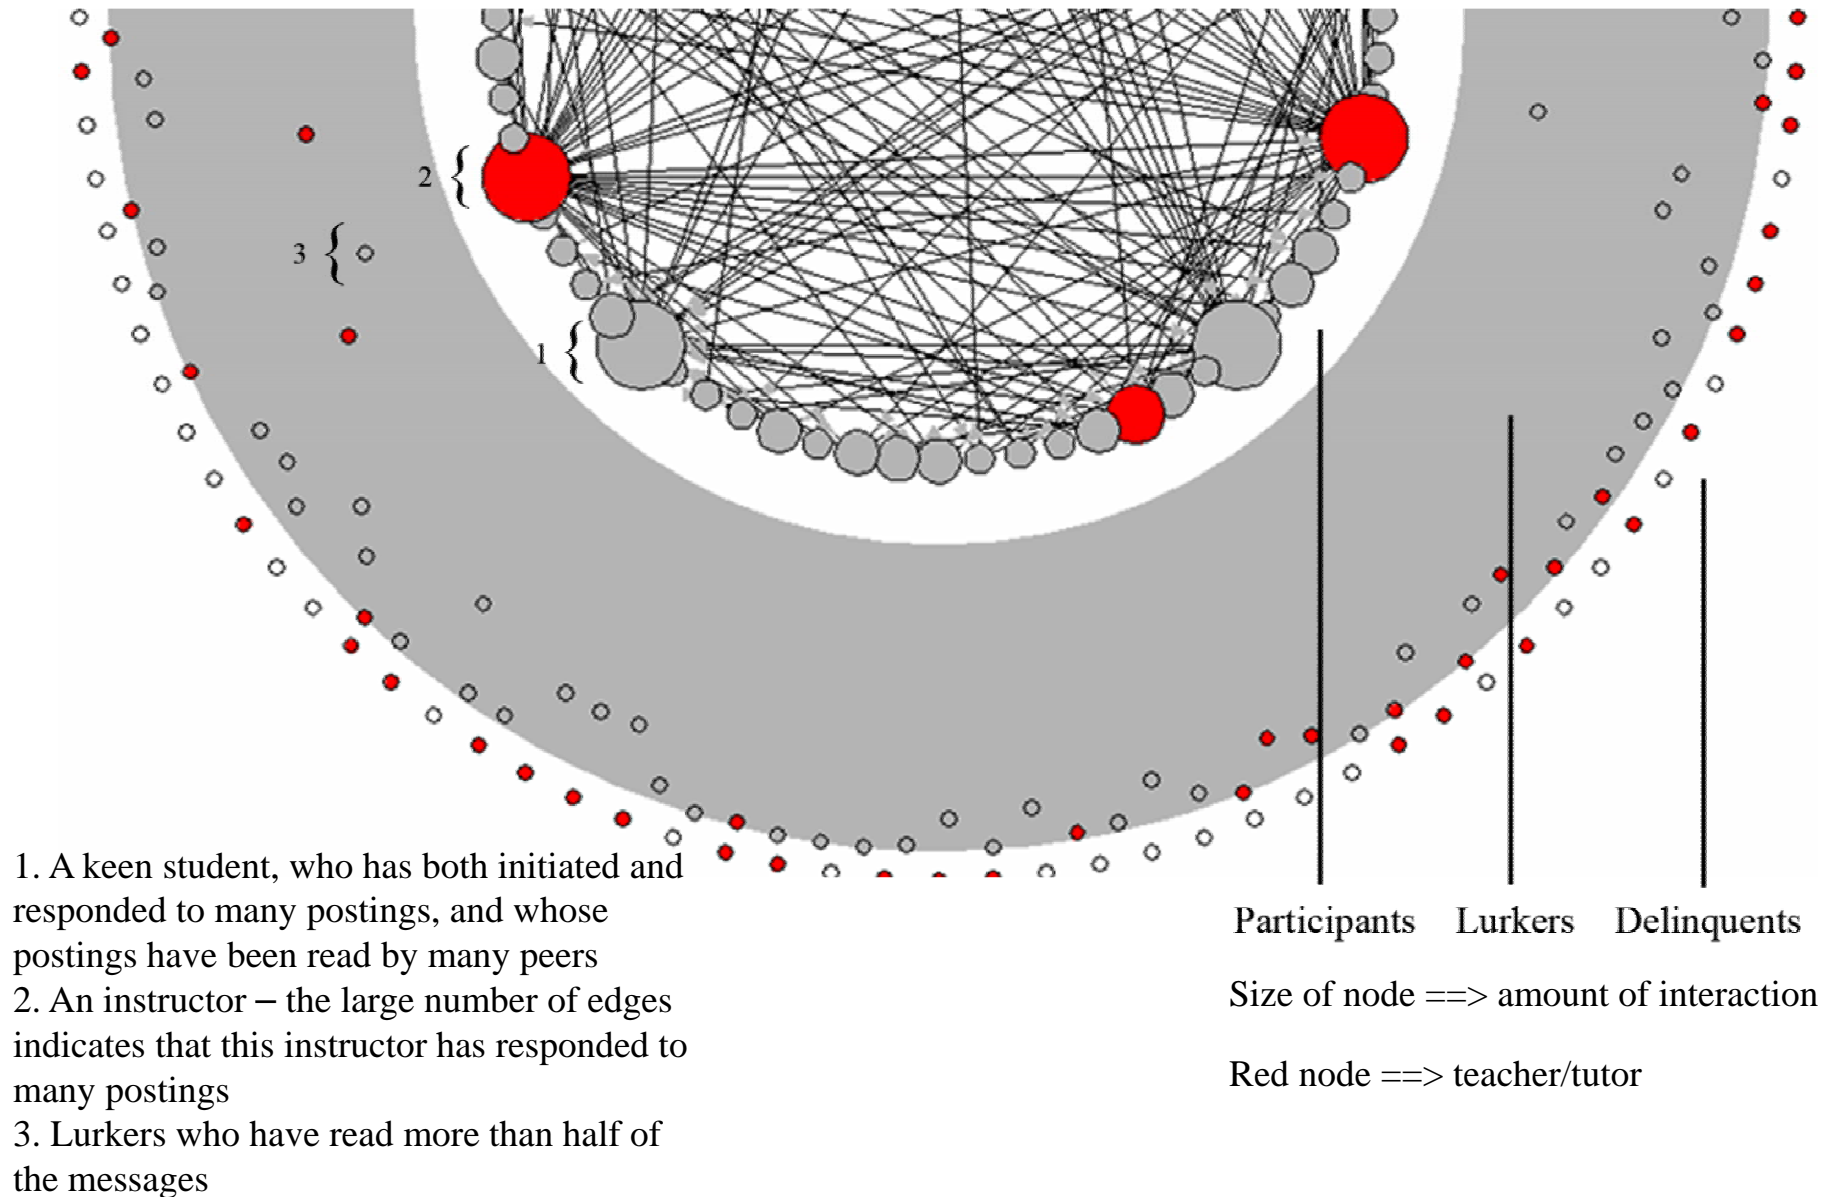
- MADMUC laboratory (Vassileva, Deters) exploring many issues of affect and community
- Also drawing on insights from iHelp
- A representative system: Comtella (2002-2008)
 - many different systems over the years
 - goal is to encourage students to participate in on-line discussions
 - various kinds of extrinsic and intrinsic rewards, most importantly reputation
 - used in a number of senior undergraduate classes over the years in various versions
 - mixed success

QuickTime™ and a
TIFF (LZW) decompressor
are needed to see this picture.


Social Network Visualization

- Visualizing social interactions in a virtual learning community - motivates participation too
 - generating *sociograms* - using iHelp discussion forum data, create a visualization of the interactions among individuals in the learning community (learners and teachers/tutors)
 - visualization can be used, for example
 - by teachers to see how individuals are performing
 - by teachers and instructional designers to see the effect of various pedagogical styles
 - to inform and motivate learners
 - Brooks, Greer, Parchoma, Daniel, Gutwin, Pinelle; CSCL-2007, CADE-2007


Sociogram Notation



Comparing Two Pedagogical Styles



A lecture oriented course with very little on-line interaction



A distance learning course with lots of on-line interaction

Privacy and Safety

Privacy

- In an approach depending so heavily on using information about learners, should we be concerned about privacy?
 - YES, but ...
 - personal agent can maintain user's privacy stance in its dealings with other agents or people
 - agents representing privacy laws and standards can also be put in the mix
 - active modelling keeps as little information around as possible, with no global picture
 - active modelling puts the focus on end-use and purpose:
 - it is not only what data the system keeps but what it can infer from that data and what it uses it for
 - these purposes will be available for inspection and execution by the user, so he or she has an idea of what the system is inferring and what it is doing with it

Privacy and Safety

- Learners need to experiment with taking on different roles in a safe environment
- Identity protection during role play
- Learners are more open and authentic with comments when identity is protected
- Identity changes complicate reputation management
 - reputation transfer from one of my identities to another
 - reputation update without giving away identity
 - accountability for bad behaviour
 - digital sanctions/imprisonment
- Anwar, Greer

The Icing



Incremental Steps to the Life-Long Learning Companion

- iHelp and our follow-up research shows that it is possible to start with a very simple learning companion and then **incrementally** add capability
 - through better user modelling
 - active modelling
 - fine-grained user tracking
 - affect
 - through reliance on surrounding social environment to scaffold weaknesses in the system
- The fully functional learning companion, as in Diamond Age is AI-complete
 - can we get there?
 - strong AI vs weak AI
 - should we get there?

Education as Social Enculturation

- Knowledge of content is always contextualized by a social environment
- Is teaching the content always of the first importance?
- Make the learner an active participant in the society
 - be it engineer, chemist, economist, psychologist, scientist, ..
- Not learning to be the Victorian lady in Diamond Age
- Social enculturation based on the society, from which derives:
 - norms of behaviour, code of conduct, body of knowledge

The Role of the Congregation

- The learner normally is surrounded by a support network to foster enculturation
 - parents, teachers, role models, mentors, peers
 - all of these are learning companions
 - do we really need an “AI companion”?
 - can we get by with a lot of resources and a collection of recommender systems?
- When the congregation is harmonious, education works seamlessly

The Power of the Educator

- There is some educational confusion when cultures collide and mixed messages are received
 - peers vs. parents
 - pop stars vs. teachers
- The learning companion could be the “voice of reason”
- The learning companion is in a highly privileged position
- Enculturation as indoctrination

The Life-Long Learning Companion is Dangerous

- Who chooses the appropriate companion?
 - would you want Nell's Primer for your daughter?
 - for your son?
- Keep the community in the picture - an inclusive and diverse society offers safety from indoctrination
- An ideal learning companion should
 - recommend resources (including human teachers, mentors and peer learners)
 - be responsible for its recommendations
 - be monitored / moderated
- And nothing more!

And/Or: The Life-Long Learning Companion is Essential

- The companion could be essential in helping Nell to navigate the vastness of cyberspace and to find appropriate communities
- The companion could be essential to protect Nell from many dangerous influences
 - how do you know what dangers your daughter or son face in cyberspace, especially from other humans lurking there?
 - the companion can protect Nell when the human social network can't, and can alert the human social network when Nell needs help
 - on duty 24/7/365

Questions, Comments, Interactions

?

Acknowledgements

- our graduate students past and present
- our colleagues in the [ARIES Laboratory](#)
- our research associates past and present
- funding from the [Natural Sciences and Engineering Research Council of Canada](#)
 - discovery grants to McCalla, Greer, Vassileva, Deters
 - LORNET networks grant to the same
- the Government of Saskatchewan through their TEL program
- private sector support: TRILabs, Parchoma Consulting Ltd.

Contact us at

mccalla@cs.usask.ca, greer@cs.usask.ca

Some References

- T.Y. Tang and G.I. McCalla, “Active, Context-Dependent, Data-Centered Techniques for E-Learning: A Case Study of a Research Paper Recommender System”, in Data Mining in E-Learning (eds. C. Romero and S. Ventura), WIT Press, Southampton, UK, 2006, pp. 97-116.
- C. Brooks and G.I. McCalla, “Towards Flexible Learning Object Metadata”, Int. J. of Continuing Engineering Education and Life Long Learning, Special Issue on the Application of Semantic Web Technologies in E-Learning (D. Dicheva, ed.), 16, 1/2, 2006, 50-63.
- G. I. McCalla, “The Ecological Approach to the Design of E-Learning Environments: Purpose-based Capture and Use of Information about Learners”. Journal of Interactive Media in Education, Special Issue on the Educational Semantic Web (eds. T. Anderson and D. Whitelock), May 2004. <http://www-jime.open.ac.uk/2004/1>
- J. Vassileva, G.I. McCalla, and J.E. Greer, “Multi-Agent Multi-User Modelling in I-Help”. User Modeling and User-Adapted Interaction J., Special Issue on User Modelling and Intelligent Agents (E. André and A. Paiva, eds.), 13 (1), 2003, 1-31.
- G.I. McCalla, “The Fragmentation of Culture, Learning, Teaching and Technology: Implications for the Artificial Intelligence in Education Research Agenda in 2010”. Special Millennium Issue on AIED in 2010, Int. J. of Artificial Intelligence in Education, 11, 2000, 177-196.