

Research Report

March 2008

Becta leading
next generation
learning

Emerging technologies for learning

Volume 3 (2008)



Contents

Foreword	03
Articles in this edition	04
Previous volumes	06
Useful resources	07
The writers	08
Growing up with Google	
What it means to education	
Diana Oblinger	11
Mobile, wireless, connected	
Information clouds and learning	
Mark van't Hooft	30
Location-based and context-aware education: prospects and perils	
Adam Greenfield	47
Emerging trends in serious games and virtual worlds	
Sara de Freitas	58
'If it quacks like a duck...'	
Developments in search technologies	
Emma Tonkin	73
Interactive displays and next-generation interfaces	
Michael Haller	91

Disclaimer

The views expressed in this publication are not necessarily those of Becta. Becta does not accept any responsibility for or otherwise endorse any information or products contained in these pages including any sources cited. Becta cannot be held responsible for any loss, damage or inconvenience caused as a result of any error or inaccuracy within these pages.

Foreword



Welcome to the third volume of Becta's *Emerging technologies for learning*. This edition builds on and is intended to complement the previous volumes from 2006 and 2007. I am sure you will agree that we have some fantastic contributions from contributors who can only be described as world leaders in their field.

Looking across the three volumes, as well as exploring a range of technologies, important themes are emerging. For example the adoption of technology in society is genuinely influencing expectations about where and how learning takes place. Educators will increasingly need to understand what these trends really mean and how to respond to related demand from learners. It will become increasingly important to understand how new technologies can enable rich, social, personalised and contextually-based interactions to support learning.

In fulfilling its role in supporting the education and skills sectors to understand the potential of emerging technologies, I'm pleased to tell you that over the next year Becta will develop new communication channels to support this. We intend not only to share research and thinking, but also to enable active participation in the building of a shared understanding of opportunities and issues. I would like to encourage you to take an active part in this process. Please keep an eye on the Becta website.

The articles in this publication should help to generate some of that discussion. I hope you find them rewarding and stimulating.

A handwritten signature in black ink that reads "Stephen Crowne". The signature is written in a cursive style and is underlined with a single horizontal stroke.

Stephen Crowne
Chief Executive



Articles in this edition

Growing up with Google: what it means to education

Diana Oblinger explores the 'net generation' who can seamlessly move between their real and digital lives. Their behaviours, preferences and expectations may be very different from those of their teachers. Diana examines the characteristics of these learners, the possibilities offered by new technologies and the skills that an education system needs to provide for the 21st century. She also argues that we need to adapt to this rapidly evolving context and goes on to explore the implications for learning space design, assessment and learning and teaching.

Mobile, wireless, connected: information clouds and learning

Mark van't Hooft looks at the implications for education of the convergence of mobile devices, pervasive wireless connectivity, and internet applications and services. These are redefining the way we live, play, work and learn, creating new opportunities for rich, personalised experiences. Mark describes some examples of how these technologies are being used to transform learning and discusses the barriers and issues to their effective use. He warns that formal education risks becoming less relevant if we don't take advantage of the new opportunities for learning enabled by these developments.

Location-based and context-aware education: prospects and perils

Location- and context-aware systems are expected to become increasingly pervasive in the near future, and here Adam Greenfield discusses some of the potential issues and pitfalls around implementation and reliance on such technologies. Adam then goes on to explore the potential of these technologies which could mean 'nothing less than a reappraisal of what we mean when we say 'education'.'

Emerging trends in serious games and virtual worlds

Sara de Freitas looks at the development of virtual worlds and 'serious games' and how we can make best use of these technologies to support better learning. They may help tackle 'the gulf between learners' experiences with technology inside and outside formal education'. She also examines the links to internet trends, web 2.0 and collaborative learning. Through discussion of the existing research and examples of the educational use of virtual worlds and serious games, Sara explores the potential of these technologies and the implications for education more widely.

'If it quacks like a duck...' Developments in search technologies

Emma Tonkin examines the problem of finding and searching digital content on the web and the limitations of current systems. She explores some of the technology developments that are beginning to address these issues, such as the semantic web, data mining, multimedia search and context-aware systems. Despite advances that offer great potential for learning through faster and more flexible access to the right information, there remains a clear role for educators and a need for better information literacy.

Interactive displays and next-generation interfaces

Michael Haller explores the potential of some emerging display and interface technologies to improve interaction with computers and facilitate collaborative activities in more natural and intuitive ways. He takes an in-depth look at interactive touch displays, drawing on both his own research and wider developments in the field. Michael also discusses the qualities of these technologies that make them appealing, and the potential benefits for the classroom.

Previous volumes

Although technology moves on incrementally as each edition is published, the articles from previous editions are forward looking enough to still be relevant now. Therefore this new volume of *Emerging technologies for learning* should be seen as complementary to the previous volumes and is not intended to replace them.

The second volume of Emerging technologies for learning (March 2007) covered six technology areas:

- Emerging trends in social software for education (Lee Bryant, Headshift)
- Learning networks in practice (Stephen Downes, NRC)
- The challenge of new digital literacies and the 'hidden curriculum' (Jo Twist, ippr)
- How to teach with technology: keeping both teachers and students comfortable in an era of exponential change (Marc Prensky)
- Games in education (Keri Facer, Futurelab and Tim Dumbleton, Becta)
- Ubiquitous computing (David Ley, Becta)

Copies can be downloaded from:

<http://www.becta.org.uk/research/reports/emergingtechnologies>

The first volume of Emerging technologies for learning (March 2006) covered five technology areas:

- Mobile learning (Geoff Stead, CTAD)
- The ambient web (Bill Sharpe, The Appliance Studio)
- Human Computer Interaction (Paul Anderson, Intelligent Content)
- Social networking (Leon Cych, Learn 4 Life)
- The broadband home (Michael Philpott, Ovum)

Copies can be downloaded or ordered from the Becta website:

<http://www.becta.org.uk/research/reports/emergingtechnologies>

Feedback

Your views and comments are valuable to us and we welcome any feedback you may have on the articles in this edition. We are also keen to hear your suggestions for topics and writers for future editions.

You can send us your feedback on the publication via email to:

emtech@becta.org.uk

Useful resources

TechNews

In order to keep up to date with relevant developments in technology we would also encourage you to sign up to Becta's TechNews service. TechNews is a free technology news and analysis service aimed at those in the education sector keen to stay informed about technology developments, trends and issues.

Each issue contains news related to the following main subject areas:

- Networking and wireless
- Multimedia
- Hardware
- Software and internet

Each subject area has a news section and a more detailed analysis piece which highlights the potential impact and likely future direction of a particular technology.

TechNews is published as a PDF once every half-term. You can subscribe to TechNews, or download it directly from the Becta website. An archive of back issues is also available.

You can find TechNews on the Becta website:

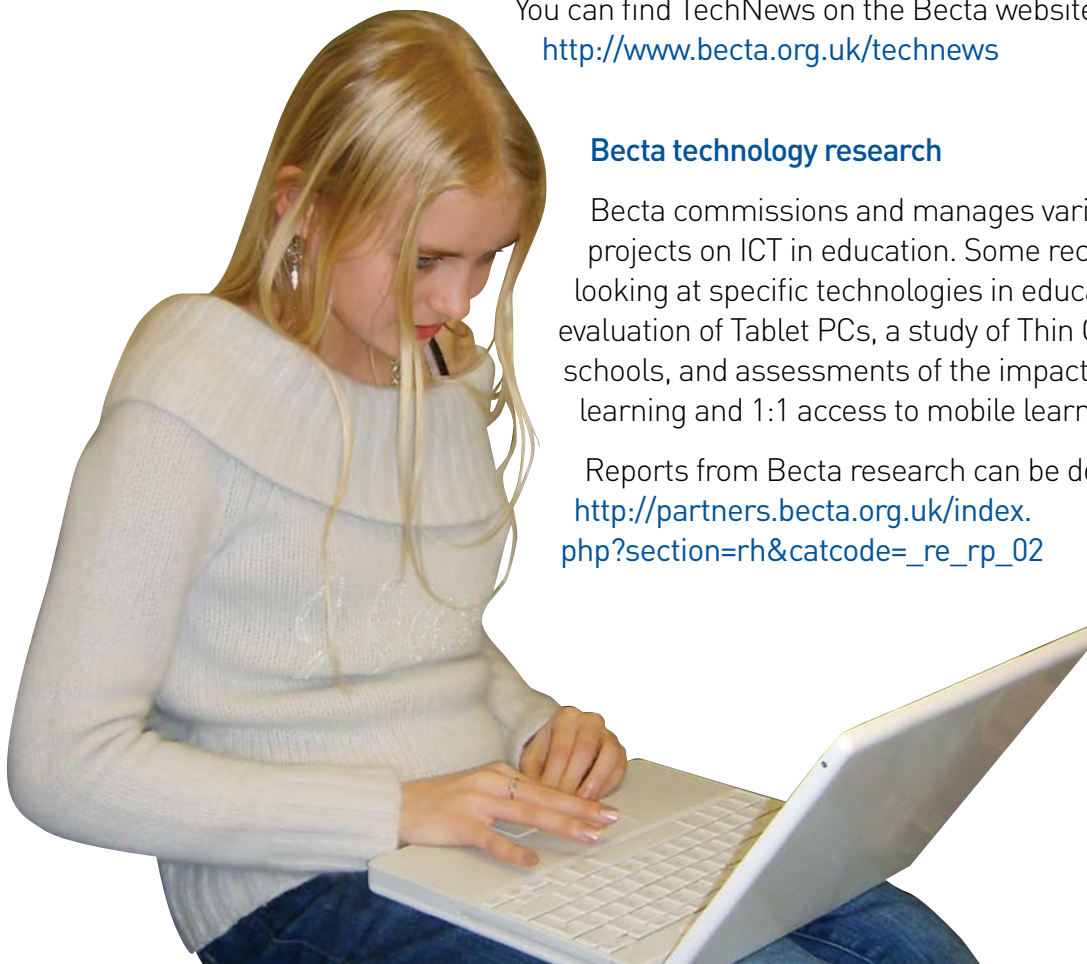
<http://www.becta.org.uk/technews>

Becta technology research

Becta commissions and manages various research projects on ICT in education. Some recent projects looking at specific technologies in education include an evaluation of Tablet PCs, a study of Thin Client technology in schools, and assessments of the impact of e-portfolios on learning and 1:1 access to mobile learning devices.

Reports from Becta research can be downloaded from:

http://partners.becta.org.uk/index.php?section=rh&catcode=_re_rp_02



The writers



Diana G. Oblinger

Dr Diana G. Oblinger is President and CEO of EDUCAUSE, an American non-profit association whose mission is to advance higher education by promoting the intelligent use of information technology. The current membership comprises over 2,200 colleges, universities and education organisations, including 200 corporations. Previously, Oblinger held positions in academia and business: Vice President for Information Resources and the Chief Information Officer for the 16-campus University of North Carolina system, Executive Director of Higher Education for Microsoft, and IBM Director of the Institute for Academic Technology. She was on the faculty at the University of Missouri-Columbia and at Michigan State University and an associate dean at the University of Missouri. She is an Adjunct Professor of Adult and Higher Education at North Carolina State University. Oblinger is internationally known for her leadership in teaching and learning with technology. She is editor or co-editor of seven books and the author or co-author of dozens of monographs and articles. Dr Oblinger has received several awards for teaching, research and distinguished service. She holds three degrees from Iowa State University and is a member of Phi Beta Kappa, Phi Kappa Phi and Sigma Xi.



Mark van't Hooft

Mark van 't Hooft, Ph.D is a researcher and technology specialist for the Research Center for Educational Technology at Kent State University and is a founding member and current chair of the Special Interest Group for Handheld Computing (SIGHC) for the International Society for Technology in Education (ISTE). His current research focus is on ubiquitous computing and the use of mobile technology in K-12 education, especially in the social studies and he occasionally writes about this on his Ubiquitous Thoughts blog [<http://ubiquitousthoughts.wordpress.com/>]. Prior to his work at RCET, Mark taught middle school and high school social studies and language arts. He holds a BA in American Studies from the Catholic University of Nijmegen, the Netherlands, and an MA in History from Southwest Texas State. He received his doctoral degree with a dual major in Curriculum and Instruction, and Evaluation and Measurement in 2005.



Adam Greenfield

Adam Greenfield is a writer, consultant and instructor at New York University's Interactive Telecommunications Program. His practice, Studies and Observations, helps clients manage challenges at the intersection of technology, design and culture, with a strong focus on issues around ubiquitous computing. His philosophy is clearly set out in his 2006 book on the subject, *Everyware: The dawning age of ubiquitous computing*. Before starting Studies and Observations, Adam was lead information architect for the Tokyo office of well-known Web consultancy Razorfish. Adam speaks frequently on issues of design, culture, technology and user experience before a wide variety of audiences. In 2007 he gave keynote presentations to the XTech conference, the Seventh International Conference on Pervasive Computing, the Monitor Group's IFA Forum, Nokia's Asia-Pacific CEO Summit, and AIGA's DUX07. He lives and works in New York City with his wife, artist Nurri Kim.



Sara de Freitas

Dr Sara de Freitas has recently taken up a new role as Director of Research at the Serious Games Institute at the University of Coventry where she leads an applied research team working closely with industry. The Institute is the first of its kind in the UK and it is envisaged that it will play a leading role in future developments of game-based learning. Formerly Sara worked as Lab Manager, Project Manager on development programmes and Senior Research Fellow at the London Knowledge Lab. Sara continues to hold a visiting senior research fellowship at the Lab.

Sara also works with the UK Joint Information Systems Committee e-Learning Development Programme in the Innovation strand, exploring the applications and developments of innovative technologies in post-16 learning. Sara's recent report *Learning in Immersive Worlds* reviews the uses of game-based learning and presents a set of case studies of practice. Sara is also working with TruSim (Blitz Games), the Vega Group PLC and the Universities of Birmingham and Sheffield on a £2 million UK Department of Trade and Industry co-funded Serious Games research and development project which will develop highly immersive learning games to address business training needs.

In 2003 Sara founded the UK Lab Group, which brings the research and development community together to create stronger links between industrial and academic research through supporting collaborative programmes and for showcasing innovative R&D solutions for the knowledge economy. Sara publishes in the areas of pedagogy and e-learning; change management and strategy development for implementing e-learning systems and educational games and electronic simulations for supporting post-16 training and learning. She also works as a consultant through her recently established partnership company: Innovatech llp.



Emma Tonkin

Emma Tonkin works as a research officer at UKOLN, which is based at the University of Bath. Following a postgraduate degree in HCI, she is currently researching a Ph.D with the Mobile and Wearable Computing group at the University of Bristol, on a topic involving the integration of human and device-level judgements of context. Her research interests include collaborative classification and social tagging, automated classification and pervasive computing. She co-moderates the DCMI Registry Community, serves as a member of the Dublin Core Advisory Board, and is a founding member of the ASIS&T SIG-TAG virtual special interest group on tagging.



Michael Haller

Michael Haller is associate professor at the department of Digital Media of the Upper Austria University of Applied Sciences (Hagenberg, Austria) and responsible for computer graphics, multimedia programming and augmented reality. He received Dipl.-Ing. (1997), Dr. techn. (2001) and Habilitation (2007) degrees from Johannes Kepler University of Linz. He is active in several research areas, including interactive computer graphics, augmented and virtual reality, and human-computer interfaces. His current focus is on innovative interaction techniques and interfaces for next-generation working environments. Currently, he leads a team of over 10 researchers and students. His research output includes journal papers and he has presented at academic conferences, and several demonstrations including ACM SIGGRAPH, Eurographics, Disney's New Technology Forum, and the Ars Electronica Festival. He has also exhibited on two occasions at the Singapore Science Center. In 2004, he received the Erwin Schroedinger fellowship award presented by the Austrian Science Fund for his time at the Human Interaction Technology Laboratory (HITLabNZ), University of Canterbury (New Zealand), and the Integrated Media Systems Center (IMSC), University of Southern California (USA).

Growing up with Google

What it means to education

Diana G. Oblinger EDUCAUSE

Describing the Net Generation

In the words of one of our students,

- My computer is the nucleus of my workspace
- When I need information I go online
- Besides IM or email my cell phone is my primary method of communication
- I'm usually juggling five things at once¹

This is the Net Generation, students who were born after 1982 – students who have never known life without the internet. Although educators may see students every day, we don't necessarily understand their habits, expectations or learning preferences. But it is obvious that technology is an integral part of their lives. To them, IM, text and Google are verbs, not applications. The Net Generation have integrated technology into everything they do, essentially putting their lives on the internet.



¹ Carie Windham. *Educating the Net Generation*. NC State University Graduate School Colloquium, September 14, 2007.

Photo: Jon Page





Photo: Department of Communications, North Carolina State University

Today's students use technology (IM, Facebook, Flickr, Skype) to be constantly connected – to friends, family, information and entertainment. Technology allows them to connect with more people, in more ways, more often. As one student explained, "Why would I call someone when I can talk to eight people at the same time on IM?"² Mobile phones, for example, aren't just for talking – they are also for texting, sending photos, accessing the Web, or watching video.³ Although older generations may find communicating via technology impersonal, the Net Generation sees technology as improving their communication. According to one student, "My parents don't understand. They think that talking online must be impersonal. Or that it leaves some sort of void. Online is how I talk. I can communicate with so many more people and manage so many more relationships. She thinks I'm more isolated than her generation – I think it's the opposite."⁴

The current generation seamlessly transitions between their 'real' and digital lives. Facebook may be the starting point for a friendship. Twitter provides constant updates on activities, feelings and observations to friends and followers. Sarah becomes Sapphire Laurasia when she enters Second Life. Students spend hours in virtual worlds and online gaming communities, where many of their friendships originate.⁵

The Net Generation demands immediate response, expecting answers at the click of a mouse. As one student commented, "They call this the information age or something, right? Look, I want things fast! I don't wait for video, I don't wait for mail, I don't wait for anything."⁶ There is an expectation of 24x7 connectivity and service. Tools like instant messaging have an inherent appeal for this generation.

Today's students bring a consumer orientation to education, which is viewed as a commodity to be consumed, acquired and accumulated.⁷ Students place a high value on the convenience technology provides, whether that means



² Carie Windham. *Father Google and Mother IM: Confessions of a Net Gen Learner*. Presented at ELI Annual Meeting, January 23, 2007. <http://connect.educause.edu/library/abstract/FatherGoogleandMothe/39228>

³ Oblinger, Diana and Marilyn Lombardi. In Press. 'Common Knowledge: Openness in Higher Education.' In Toru Iiyoshi and M. S. Vijay Kumar (Eds), *Opening Up Education: The Collective Advancement of Education through Open Technology, Open Content, and Open Knowledge*. MIT Press, Cambridge MA.

⁴ Carie Windham. *Educating the Net Generation*. NC State University Graduate School Colloquium, September 14, 2007.

⁵ Julie Evans. *K-12 Students Speak Up about Technology: Are We Listening?* Presented at ELI Annual Meeting, January 22, 2007. <http://connect.educause.edu/library/abstract/K12StudentsSpeakUpAb/39232>

⁶ Carie Windham. *Educating the Net Generation*. NC State University Graduate School Colloquium, September 14, 2007.

⁷ Mark Taylor (2006). 'Generation NeXt Comes to College: Today's postmodern student.' <http://globalcscce.edu/tirc/blog/files/Gen%20NeXt%20handout%2006%200ln.pdf> Accessed September 19, 2007.

accessing course material from anywhere at any time or being able to see course grades as soon as they are posted. Many students describe education as a business where efficient, convenient, technology-mediated transactions are expected.⁸ Consumerism can sometimes bleed into an entitlement culture, however. An increasing number of students – and their parents – expect academic success with little academic effort.⁹

Their learning styles are influenced by the immediacy and visual richness of the environment they have grown up in, particularly television and the internet. Net Generation students expect to be engaged by their environment, with participatory, sensory-rich, experiential activities (either physical or virtual) and opportunities for input. They are more oriented to visual media than previous generations – and prefer to learn by doing rather than by telling or reading. “Don’t just tell us – let us discover.”¹⁰ To illustrate, one student described how she learned about video. “Well...I opened up the camera box, started messing around, and then figured out how to upload it. Took a while. Had to Google it a few times to figure out how to splice stuff together. Just took an hour or so.”¹¹ They teach themselves how to use technology – or learn it from peers.

Harbingers of change

One reason for trying to understand the Net Generation is that our students may be harbingers of change. Their habits, expectations and behaviours may anticipate what the rest of society will come to consider as its culture or norms. In fact, indicators suggest that society’s shared beliefs, values, customs and behaviours are being reshaped by globalisation and technology. These changes apply across the spectrum of age and occupation – not just to young people.

Communicating and connecting

The internet is a major channel for socialising. For students of all ages, communication drives many of their uses of IT. In a survey involving almost 800,000 school-aged children in the US, 23 per cent say they are connecting with people around the country – not just in their class, or their neighbourhood.



⁸ A Net Gen Student Perspective: Technology in Higher Education. May 7, 2007. <http://www.educause.edu/ELIWEB075>

⁹ Mark Taylor. 2006. ‘Generation NeXt Comes to College: Today’s postmodern student.’ <http://globalcscc.edu/tirc/blog/files/Gen%20NeXt%20handout%2006%2001n.pdf> Accessed September 19, 2007.

¹⁰ Carie Windham. Father Google and Mother IM: Confessions of a Net Gen Learner. Presented at ELI Annual Meeting, January 23, 2007. <http://connect.educause.edu/library/abstract/FatherGoogleandMothe/39228>

¹¹ Carie Windham. Educating the Net Generation. NC State University Graduate School Colloquium, September 14, 2007.



Photo: Department of Communications, North Carolina State University

Seventeen per cent are connecting with young people around the world. Approximately one third of their friends are people they've never met face-to-face. 'Globally, the average young person connected to digital technology has 94 phone numbers in his or her mobile, 78 people on a messenger buddy list, and 86 people in their social networking community.'¹²

Social networks, such as MySpace, Facebook and Bebo, are part of how students communicate. From Facebook they send messages (email or instant messages), they blog, they post pictures, and they 'poke' each other. But communicating and connecting isn't limited to text. Students share photographs through sites like Flickr or by emailing photos from mobile phones. An increasing number are using geolocation services to find friends who are in the area.

Moreover, communicating and connecting are not limited to the real world or real people. An increasing number of students meet friends through online gaming communities, often establishing friendships with students from other continents – something that would be impossible in the real world. Others establish connections in virtual worlds, such as Second Life, in which users are represented by avatars and by pseudonyms, which are also common among bloggers. For users of online worlds, identity has become a flexible concept – being 'yourself' is not necessarily limited to your physical being or given name. Pseudonyms and avatars can expand one's identity.

No matter what medium, communication is an imperative for the Net Generation. Among students surveyed, 100 per cent communicate with others each time they log onto the network; 70 per cent check IM as the first thing they do when they turn on their computer.¹³

Collaborating and co-creating

A host of Web 2.0 technologies enable collaboration and co-creation activities, perhaps exemplified by Wikipedia, in which users write and edit their own encyclopaedia. Since the Web has become our information universe, we have developed a do-it-yourself approach to finding information. Whether it is booking travel, researching a paper, or seeking entertainment, the first stop is likely to be the Web.

The Web is no longer just a way to receive information – it is a medium for commenting, collaborating and creating. Blogging, where anyone can create



¹² MTV, Nickelodeon and Microsoft challenges assumption about relationship between kids and Digital Technology. July 8, 2007. http://advertising.microsoft.com/sverige/NewsAndEvents/PressRelease.aspx?Adv_PressReleaseID=630

¹³ MTV, Nickelodeon and Microsoft challenges assumption about relationship between kids and Digital Technology. July 8, 2007. Ibid.

and publish their ideas and opinions, is one form of internet-based self-expression. More than 50 million blogs were created by mid-2006; estimates indicate that two new blogs are created each second.¹⁴ The power of blogs goes beyond self-expression. For example, the 2007 uprising in Burma was made known to the outside world through blogs. Blogs are influencing mainstream media through the emergence of 'citizen journalism'. Individuals can make their voices heard worldwide, whether through blogging, podcasts, or sharing opinions on sites like Digg.com. Due to the creation and co-creation possible on the Web, control of information has shifted from being highly centralised (through, for example, major media outlets) to highly distributed.

Collaboration and co-creation enable 'collective intelligence' or distributed cognition. For example, 'We Are Smarter Than Me' is a book project where the material was developed using a wiki. Collective intelligence recognises that nobody knows everything but that everyone knows something. Diverse groups of people can pool knowledge, research, debate, and create new insights thanks to a networked culture that is redefining power structures.¹⁵

Collective intelligence is a powerful force that is reshaping what our traditional IT infrastructure is capable of. One of the distinctions between traditional IT infrastructure and cyber-infrastructure (or e-science) is the enablement of distributed cognition, where the infrastructure has a role in creating social connections and facilitating the work of virtual organisations. 'Professional scientists and amateur enthusiasts form virtual research communities advance the study of astronomy, ornithology, and other fields that rely on the collection of large data sets. No longer the exclusive purview of credentialed scientists, data collection and analysis is open to all interested parties. With distributed cognition, contributors come from all walks of life, information flows in multiple directions, and a bottom-up energy drives discovery.'¹⁶

Moving beyond text

Students seem to be more at home with images (icons, video, photos) than text – the opposite of what most educators consider their comfort zone. According to a 2005 study, more than one half of all American teens – and 57 per cent of teens worldwide who use the internet – could be considered media creators, producing blogs and Web pages, posting original artwork, stories, or videos



¹⁴ Amanda Lenhart and Suzannah Fox (2006). 'Bloggers: A portrait of the internet's new storytellers.' <http://www.pewinternet.org/pdfs/PIP%20Bloggers%20Report%20July%2019%202006.pdf>

¹⁵ Henry Jenkins. November 27, 2006. Collective Intelligence vs. The Wisdom of Crowds. http://www.henryjenkins.org/2006/11/collective_intelligence_vs_the.html. Accessed October 2, 2007.

¹⁶ Diana Oblinger and Marilyn Lombardi (2007). 'Common Knowledge: Openness in Higher Education.' In Toru Iiyoshi and M. S. Vijay Kumar (Eds), *Opening Up Education: The Collective Advancement of Education through Open Technology, Open Content, and Open Knowledge*. MIT Press, Cambridge MA.

online or remixing the online content of others.¹⁷ Students capture images with mobiles or video cameras then share experiences with friends and strangers alike through MySpace or Facebook-like environments.¹⁸ Socialising involves images as well as text. And, visual literacy is an important part of digital literacy.

Today's visual options are expanding rapidly. Students email photos from their mobile phones but also post images on Flickr. Video can be shared on YouTube or on sites designed for young people such as UthTv. But images represent only one possible dimension. For example, Flickr goes beyond just photo sharing – users can geotag photos, pinpointing the photo's exact latitude and longitude. Images can be integrated with Google Maps as well, allowing users to populate locations with their own tags and documentation, sharing a personal history associated with the space.¹⁹

Another media form is the mashup, which combines stand-alone technologies into a novel application, allowing users to put together different types of data. Mapping mashups, in which maps are overlaid with information, may be the best known example of this rapidly growing genre. Some mashups provide details for specific locations. Others bring together different data sources, such as combining crime data with location information (ChicagoCrime.org for example). A music mashup mixes tracks from two different source songs. Tools (for example, Google's Mashup Editor, Intel's mashup maker, or MIT's Piggy Bank [http://simile.mit.edu/wiki/Piggy_Bank]) allow non-technical individuals to mix up data, find new meaning, and present it in interesting ways. Educationally, mashups can be extremely valuable (helping students integrate previously disparate types of information), but they are not without their cautions (such as use of others' intellectual property).

Many educators express concern that students do not read – well or enough. Yet we find ourselves in an increasingly visual world. Graphic representations sometimes reveal relationships in data, for example, that could not be readily discerned from tables of numbers. Applications are emerging that capitalise on visual relationships (for example, Quintura [<http://www.quintura.com>] which displays visual relationships; their website promotes Quintura for kids as the first 'visual search engine for the youngest Web users'). Visuwords [<http://www.visuwords.com>] allows users to look up words to find their meaning as well as discover associations with other words and concepts, which are displayed graphically.



¹⁷ Henry Jenkins with Katie Clinton, Ravi Purushotma, Alice J. Robinson, and Margaret Weigel (2006). 'Confronting the Challenge of Participatory Culture: Media Education for the 21st Century.' http://digitallearning.macfound.org/atf/cf/%7B7E45C7E0-A3E0-4B89-AC9C-E807E1B0AE4E%7D/JENKINS_WHITE_PAPER.PDF Accessed February 2, 2007.

¹⁸ Cyprien Lomas and Diana Oblinger (2006). Student Practices and their Impact on Learning Spaces. In *Learning Spaces*.

¹⁹ Ibid.

But even visual images may merit enhancement. A team of researchers at Zhejiang University is developing the Emoplayer, a tool that highlights the emotional state of the characters on a video timeline. The emotional markers make it easier for users to navigate video content compared to standard media players. Users who search with emotional markers find the desired segments more quickly than those using traditional search techniques.²⁰

Blending real and virtual worlds

Today's students grow up playing video/computer games, join guilds in massively multiuser online games (such as World of Warcraft), and participate in virtual worlds (such as Second Life). Students blend the real and virtual worlds, moving seamlessly between them.

For example, immersive multiplayer virtual environments let players participate in new worlds, inhabiting roles that would otherwise be inaccessible to them. They allow people to think, act, and talk in new ways. Rather than relying on words and symbols, learners experience the virtual world – participating in a mission to Mars, experiencing a native culture, travelling through the human body. In these virtual worlds students assume the role of an expert, experiencing how a particular discipline thinks about and solves problems, as a physicist, an astronaut, an anthropologist, or a physician. As a member of a community of novices and experts, learners develop knowledge, skills and values; novices are exposed to the ways professionals deal with problems, mirroring the practice of being an expert.²¹

Augmented reality is another example of blending real and virtual environments, in which digital information is superimposed on the real world, many times involving a handheld device or mobile phone. Users can view the Empire State Building, for example, and superimpose on its image the names of businesses in the building, the cost of visiting the building's observatories, or hours and menus for its five restaurants.²² Such techniques can be used for nature walks, museum tours, or simulations (Environmental Detectives for example²³).

Students are also taking advantage of geotagging to leave 'virtual messages' that can be read by handhelds at specific physical locations. Some campuses are encouraging students to geotag places of importance to them or to explain the history of the campus. Others use geotagging as the basis of games, such as virtual scavenger hunts.



²⁰ <http://www.newscientist.com/blog/technology/2007/07/adding-emotions-to-online-video.html>

²¹ Diana Oblinger (2006). 'Games and Learning.' EDUCAUSE Quarterly 29(3): 5-7.
<http://www.educause.edu/ir/library/pdf/eqm0630.pdf>

²² <http://www.technologyreview.com/Infotech/18291> Accessed October 2, 2007.

²³ <http://education.mit.edu/ar/ed.html>

Misconceptions/assumptions

Technology preferences

Although the effortless use of technology by this generation may be striking, appearances are often deceptive. While this generation shows no fear of technology, 'digital comfort' does not necessarily mean technology proficiency – particularly with academic tools. When queried, students often advise, "Don't assume we can plug a formula into Excel. Or that we know how a wiki works. Sometimes it is just new to us."²⁴ Nor does comfort with technology equate to a full appreciation of issues such as intellectual property, privacy or security. When asked, most students confess, "Sometimes we just don't think about what we're doing online."

Information fluency

Students need to know how to find and use information, and technology is a critical enabler. However, only 31 per cent of information searches are successful.²⁵ Just because students know how to open a Web browser, educators should not assume that everyone knows how to search for information. And just because students can find information doesn't mean it is reputable or truthful. As a do-it-yourself culture where we find information for ourselves online, we need the skills to vet what we find, understand the context in which the information is situated, and adjust our interpretation accordingly. Also, in a cut-and-paste and mashup-friendly environment, students must develop an appreciation for intellectual property and the work of others.²⁶

Access to technology

In a wired world it is easy to assume that all students have access to a computer and the network, whether at home or at school. However, a digital divide still exists in many communities – one defined as more than just having access to a computer. A 'second-level digital divide' may exist based on machine vintage, connectivity, online skills, autonomy and freedom of access, and computer-use support.²⁷ Another common assumption is that all students are attracted to technology. No group is entirely homogeneous. Not all students have computers, not all are skilled users, and not all want to use technology.



²⁴ Carie Windham. Father Google and Mother IM: Confessions of a Net Gen Learner. Presented at ELI Annual Meeting, January 23, 2007. <http://connect.educause.edu/library/abstract/FatherGoogleandMothe/39228>

²⁵ Study from Information Online conference.

²⁶ Diana Oblinger (2007). 'Becoming Net Savvy.' EDUCAUSE Quarterly 30(3): 11–13. <http://www.educause.edu/ir/library/pdf/eqm0731.pdf>

²⁷ Hawkins, Brian and Diana Oblinger. July/August 2006. 'The Myth about the Digital Divide.' EDUCAUSE Review 41(4): 12–13. <http://www.educause.edu/ir/library/pdf/erm0647.pdf>

Maturity

It is easy to assume that learners – with their tech-savvy attitudes and world-wise veneer – have greater maturity than their years. We are cautioned that this presumption of maturity is unfounded on many levels. In a multi-tasking, fast-forward world, learners may not be stopping to reflect on what they know, how they behave, and the values they hold. In fact, the tendency of young people to not be reflective – to pause, think, and ponder – may simply be a characteristic of youth. However, in an environment where students are posting their lives on the internet, stopping to consider what they are doing and its future impact is essential. It is all too easy for students to follow their peers and not stop to consider whether their behaviour mirrors their values and beliefs. Just as students may not reflect on their online behaviour, they may not reflect on their learning habits.

Becoming net savvy

Although we may assume students are technologically savvy, most indications are to the contrary. As a result, educators, parents and communities must develop the policies and practices that students need. Although specific policies must be determined in a local context, some questions may guide their development. For example:

- Whose responsibility is it to help students be thoughtful in their use of the internet, whether that means limiting the amount of personal information posted online or not assuming that ‘childish pranks’ are easily erasable on the internet?
- In an environment in which it is all too easy to cut-copy-and-paste, how do we ensure students develop respect for intellectual property and the ethical use of information?
- How do we ensure that all students have the opportunity to develop the requisite technology skills, whether or not they have access to a computer at home?
- How do we ensure that information fluency becomes a habit of mind rather than of an isolated library requirement if parents, teachers, and staff do not integrate into their daily interactions with students?
- How do we ensure that students develop the critical thinking skills necessary to survive and thrive in an age when anything (true or false) can be found on the internet? What programs will help students develop understanding based on evidence, critical thinking, values, and dialogue rather than the first item on a Google search?

The goal is to ensure that our students – and the rest of us – are net savvy. 'Becoming net savvy isn't a one-time affair – it is a lifelong educational process – and something that should be integrated into all aspects of our lives. Ensuring that we are all net savvy will require a team effort – libraries, IT, instructors, parents, community centres, and others. It calls for a protracted effort, starting in the early years, extending throughout life. Being net savvy – or not – is no longer an option. It is an imperative in the age of information.'²⁸

Implications for education

Student needs and expectations, the technological and pedagogical tools available, as well as what it means to be educated in the 21st century are leading educators to envision education that is interactive, engaging and challenging. We are also learning that students may be among the best advisors on how to strengthen education.

What it means to be educated

Education designed for the economic processes and institutions of the 20th century may only prepare students to work in organisations that are now rapidly becoming obsolete. In the future, more students will run their own businesses rather than work for others. More jobs will be created in small, medium, and entrepreneurial firms than in large multi-nationals. Workers must constantly, quickly, and efficiently learn new skills and information. Today's students must graduate able to deal with ambiguity and capable of higher-order analysis and complex communication.²⁹

Learners need skills that go far beyond reading, memorisation and communication. Educational institutions have an obligation to help students cultivate those skills that learners have the most difficulty attaining on their own, such as:

- 1 judgement, or the ability to distinguish the reliable from unreliable information
- 2 synthesis, or the capacity to follow the longer argument or narrative across multiple modalities



²⁸ Diana Oblinger (2007). 'Becoming Net Savvy.' EDUCAUSE Quarterly 30(3): 11–13. <http://www.educause.edu/ir/library/pdf/eqm0731.pdf>

²⁹ Dede, Chris, Spence Korte, Robert Nelson, Gil Valdez, and David J. Ward (2005). 'Transforming Learning for the 21st Century: An Economic Imperative.' Naperville, IL: Learning Point Associates. <http://www.learningpt.org/tech/transforming.pdf> Accessed February 2, 2007. Page 3.

- 3 research, or the activity of searching, discovering, and disseminating relevant information in a credible manner
- 4 practice, or the opportunity to learn-by-doing within authentic disciplinary communities
- 5 negotiation, or the flexibility to work across disciplinary and cultural boundaries to generate innovative, alternative solutions.³⁰

Expansion of learning opportunities

Today's students are motivated by solving real-world problems, preferring to do rather than simply listen, and most educators consider learning-by-doing the most effective way to learn. There is a significant difference between learning about physics and learning to be a physicist, for example. Isolated facts and formulae do not take on meaning and relevance until learners discover what these tools can do for them. Immersive and authentic learning environments, such as simulations, visualisations, haptics, augmented reality or virtual worlds can be both engaging and motivating.³¹

Virtual worlds

A virtual world is an online environment whose 'residents' are avatars representing individuals participating online. The functioning of a virtual world can mirror that of the real world, or it can allow residents to do such things as fly, wander around underwater, or teleport themselves to other locations. Today's virtual worlds are immersive, animated, 3D environments that operate over the internet, giving access to anyone on the world. Dartmouth College, for example, is creating a virtual world to train community emergency response teams. Harvard University created River City, a virtual world that presents users with an outbreak of disease, allowing them to move through the environment, make inquiries, and examine data to try to discover the source of the illness.³²

Consider, for example, a virtual world that is a clinic in which Marie, a first-year medical student, assumes the role of a doctor through her avatar. The world closely mimics a real doctor's office, and Marie 'walks' her avatar into an examination room, where she finds a patient waiting. By typing dialogue, Marie speaks to the patient, introducing herself, and the patient stands up and introduces himself. An important element of becoming a doctor is learning how to interact well with patients, and the virtual clinic gives Marie an opportunity to do that. Marie's avatar and the patient – who is the avatar of



³⁰ Jenkins, et. al. (2006), 4.

³¹ Marilyn Lombardi (2007). Authentic Learning for the 21st Century: An Overview. <http://www.educause.edu/ir/library/pdf/ELI3009.pdf>

³² 7 Things You Should Know About Virtual Worlds. June 2006.

another medical student – talk about what’s bothering the man, which turns out to be a flare-up of gout. Controlling her avatar through the interface, Marie asks questions, reads the patient’s chart, and conducts a clinical examination. Faculty monitor the sessions in the virtual world and can interrupt to offer suggestions about better ways to ask certain questions and to earn the patient’s confidence. Because it has both real and simulated components, Marie finds that the interface to the virtual world requires her to be extremely conscious of her actions and the words she uses when interacting with patients, a skill she understands she will need to apply when in a real office with real patients.³³

Remote instrumentation

Remote instrumentation provides control of scientific instruments over a network from remote locations, such as telescopes, shake tables, or electronics equipment. Because of their expense and complexity, many specialised scientific instruments are out of the reach of some institutions; even for those that can afford them, scheduling and other logistical issues may prevent full utilisation of those tools. Remote access addresses issues of access and efficiency, providing students with real experiences, ultimately improving educational quality.³⁴

Think of a scenario in which a consortium of universities and oceanographic institutes installs a bank of sensors and testing equipment on the ocean floor, as well as atmospheric instruments at the water’s surface. The project includes underwater cameras, water-sampling and analytical tools, and other devices to monitor and measure a wide range of oceanographic and meteorological activity. The equipment is linked to the institutions in the consortium, where researchers and faculty can control the devices. Dr Morgan, a biology professor at one of the participating institutions, frequently turns to the undersea lab to demonstrate experiments and show students the results. He can control video cameras on the sea floor, collect water samples and feed them into a tool that analyses pH and other parameters, and even position the water sample under a microscope, all through a browser interface. He uses data from the instruments to show the class correlations between atmospheric conditions and the conditions of the water and marine life. Outside class, students can perform most of the same manipulations from their dorm rooms, repeating experiments to see how the results change over time or devising new tests. Students at other universities can also access the instruments, and the undersea lab is in use most hours of the day, transmitting data and observations across the internet.³⁵



³³ Ibid.

³⁴ 7 Things You Should Know About Remote Instrumentation. April 2006.

³⁵ 7 Things You Should Know About Remote Instrumentation. April 2006.

Augmented reality

Many of today's students move seamlessly between the physical and virtual worlds, and blending the two can provide a valuable learning environment. Augmented reality does this, but adding digital information to a real object or place. Unlike virtual reality, augmented reality does not create a simulation of reality but takes a real object or space and incorporates technologies that add contextual data to deepen a person's understanding of the subject.³⁶

Josie, a student who missed a field trip to a botanical garden, for example, might use augmented reality make up that trip on her own, using a guided tour that the professor created. She would go to the garden and launch the tour, which is loaded on her PDA. Josie starts at a cedar, which, according to the material on her PDA, is more than 500 years old. As she approaches the tree, the GPS in her PDA notes her location, and she hears a recording of the professor giving his theory about the role that trees like this one play in the ecosystem. As she moves through the garden, she selects photos and movies of other trees, depicting the history of the garden, seasonal differences, and changes that have occurred. In addition to the material supplied by the professor, Josie can also download the notes, photos, and keywords that the rest of her class recorded when they took the field trip. She adds her own observations, assigning appropriate keywords and GPS coordinates to her notes.³⁷

Mapping mashups

Mapping mashups overlay data on maps with clickable markers showing specific points of interest. Data interoperates with an online mapping service, putting information in a geographic context. In a wide range of academic disciplines, understanding the geographical context of places and events is central to a deep comprehension of the subject matter. Mapping mashups do this by combining a mapping tool with other applications and online resources to create interactive learning experiences.

Consider a history course about World War II, for example. The instructor, Dr Martinez, develops a mapping mashup that represents major events leading up to and during the war. Users can navigate around the world with the tool, zooming in and out, showing the map with dates and events superimposed. When students zoom in on Europe, for instance, they see markers scattered around the continent and into Asia. When clicked, each marker opens a pop-up box that names the location, explains what happened there and when, and shows a photograph of that site. The markers correspond to important battles, political events, treaties that were signed, and cities such as Vichy, the wartime capital of France. The text in each box also includes links to articles that talk in depth about what took place in each location and its significance.³⁸



³⁶ 7 Things You Should Know about Augmented Reality. September 2005.

³⁷ Ibid.

³⁸ 7 Things You Should Know about Mapping Mashups. July 2006.




Photo: Department of Communications, North Carolina State University

Data visualisation

Data visualisation is the graphical representation of information. Bar charts, scatter graphs, and maps are examples of simple data visualisations that have been used for decades. Although data visualisations have long been used in academic settings, instructors are using new technologies that combine the principles of data visualisation with powerful applications and large data sets. The results are rich, compelling visualisations, including sophisticated images as well as animations, that help students understand concepts more quickly and deeply than with older tools.³⁹

Imagine a data visualisation tool designed to help athletes improve their performance. Olivia, a graduate student in kinesiology, works with faculty in the computer science department to develop a data visualisation tool that correlates data including average, resting, and maximum heart rate; lactate threshold, blood oxygenation, and other variables. The tool creates 'pictures' of the performance of members of the track team. Each image looks something like a 3D map of hilly terrain, with colour differences that reinforce the contours. Peaks represent efficient performance, and the runners can see in visual terms how their performance is affected, for example, by exceeding their aerobic thresholds. During treadmill workouts, the athletes can watch a computer screen that displays visualisations based on data collected in real time. In this way, the runners can watch a representation of their efficiency and see how it changes as they modulate their effort or change their breathing patterns. Over the course of a semester, Olivia could use the tool to help the runners understand – through the visual representations of effort and efficiency – the factors that most benefit and hurt performance. One runner might discover that her performance is maximised by a very even effort, whereas another might find that her best time comes from varying her intensity. The visualisations might show previously hidden correlations between weight training and aerobic capacity, leading to changes in the training programme for the whole team.⁴⁰

Digital and convenient

Today's students use the computer as their notebook, locker, backpack, and organiser.⁴¹ They expect technology to provide solutions for their wants and needs. Students say they want more 'learning-on-the-go' options and mobile device services to align with their mobile lifestyle. Others ask that applications be integrated so students can access their schedules, campus events, and other information from the same login.



³⁹ 7 Things You Should Know about Data Visualization. October 2007.

⁴⁰ Ibid.

⁴¹ D. Levin and S Arefeh. August 14, 2007. The Widening Gap Between Internet Savvy Students and their Schools. Pew Internet & American Life Project.

Students expect cross-platform access to content, the ability to download and upload material, and the integration of digital media in their learning tasks.⁴² They ask that course content, class notes, lectures and syllabi be searchable with common tools such as Google and available 24x7. Students also suggest that a wide array of courses should be available online, providing greater flexibility than traditional class schedules and that lectures be available as video-on-demand.⁴³

The ways institutions communicate with students today, mostly in text, are described by students as 'flat'. Students suggest more visual options. They see opportunities to use multimedia to enrich services as well as courses. For example, students suggest that information in a degree audit would be more understandable with a graphic interface rather than lines of text. They advocate maps pinpointing open parking spaces or open computers in the library.

Students also suggest that institutions might do more to foster a sense of community among students. Remembering that our current generation of learners does not limit the definition of communication to face-to-face interaction, suggestions include integrating social technologies in institution websites, allowing students to share photos, using social bookmarking, and blogging.

Peer production

The Economist has declared that the era of peer production has arrived.⁴⁴ 'From Amazon.com (where much of the value comes from millions of customer reviews) to MySpace to Craigslist, the most successful Web companies are building business models based on user-generated content.'⁴⁵ It is a model that has instant credibility and applicability in education. The open-source movement epitomises peer production.

Open educational resources, or the sharing of teaching materials (content modules, courseware, learning objects, online learning communities), is an increasingly popular model. Notable examples include MIT's OpenCourseWare or Open University's OpenLearn. While an open source model brings advantages, it challenges our existing educational practice of assessment and attribution. How does the community judge the quality of its collective output? Who receives credit for creating it?



⁴² Ibid.

⁴³ A Net Gen Student Perspective: Technology in Higher Education. May 7, 2007. <http://www.educause.edu/ELIWEB075>

⁴⁴ *The Economist*. April 22, 2006.

⁴⁵ Brenda Gourley. March 23, 2007. 'How Technology is Shaping Educational Agendas' Presented at the 36th Scottish Council Forum.

Control and authority

In an environment where an instructor's statements can be immediately verified or discounted – based on information on the internet – and in a world of collective intelligence, 'control' and 'authority' take on different meanings. Whether due to the internet or the current generation's empowerment, the automatic acceptance of academic authority is a thing of the past.⁴⁶

Institutions cannot control content because so much is on the Web, and many web-based resources have been created by amateurs rather than by institutionally appointed authorities. In addition, in a Web 2.0 world, those resources might have been remixed by someone else. The traditional assumption that information comes from the library, which implies quality control by publishers, peer-review panels, and librarians, is no longer guaranteed to be true. Students seek information on the internet first; visiting the library stacks may be a relatively rare occurrence. In fact, many students have a very limited working knowledge of the library. What has become increasingly important is helping students develop information fluency or becoming 'net savvy'.⁴⁷

Authority in an internet world may be based more on reputation and the strength of one's network than on educational credentials or position. Where in one's academic credentials do you cite your reputation as a blogger or tout how popular your YouTube video is?



⁴⁶ Mark Taylor (2006). 'Generation NeXt Comes to College: Today's postmodern student.' <http://globalcscce.edu/tirc/blog/files/Gen%20NeXt%20handout%2006%200ln.pdf> Accessed September 19, 2007.

⁴⁷ <http://www.educause.edu/ir/library/pdf/ELI3006.pdf>, <http://www.educause.edu/ir/library/pdf/ELI3008.pdf>



Photo: Department of Communications, North Carolina State University

Newspapers and media outlets have learned that citizens may have the hottest stories or the most relevant photos that can augment the historic strength of media channels. Wikiversity 'put [traditional institutions] on notice that even the world of education will be challenged by this new world where lectures are turned into conversations among people formerly known as the audience.'⁴⁸

Although today's students don't necessarily generate course content, there is nothing to preclude that in the future. Some envision a global networked community of tutor and student volunteers who would produce resources that are evaluated and ranked by the community as a whole. The model might mirror today's open source or internet 'reputation' model. Rather than the content being fixed, it could be distributed globally to be sampled, mashed up, remixed, and re-contextualised for local use.⁴⁹

Learning spaces

Space can open opportunities for new pedagogies, interactions, and connections, particularly since wireless technology makes it possible for almost any place to be a learning space. Beyond classrooms, institutions are redesigning space to ensure that student learning and interaction with faculty can take place across the entire campus. Libraries are being transformed to information commons, where the floor-to-ceiling book stacks and carrels designed to ensure solitude and silence are giving way to open spaces where technology is integrated with talk and food. Whether the conversation is about math, music, or football, space can bring people together providing them with opportunities to learn from others – academically, socially, culturally.⁵⁰ Institutions are:

- designing space around learning rather than instruction
- creating socially catalytic spaces – places where people meet, congregate, and socialise
- integrating technology in spaces and putting services where students are
- involving users in the design of spaces.



⁴⁸ Brenda Gourley. March 23, 2007. 'How Technology is Shaping Educational Agendas' Presented at the 36th Scottish Council Forum.

⁴⁹ Hylén, Jan (2006). 'Open Educational Resources: Opportunities and Challenges.' <http://www.oecd.org/dataoecd/5/47/37351085.pdf> Accessed February 13, 2007

⁵⁰ Diana Oblinger (2007). What Growing up with Google may mean to graduate education. Council of Graduate Schools Communicator.

Changing mental models

An ongoing challenge is changing educators' mental models and assumptions about learners, technology, and the skills needed in a global economy. The default is a model in which learning and teaching take place in a classroom, yet we know that most learning occurs outside formal educational settings. Even those who go beyond the limitations of assuming learning happens in the classroom may have difficulty getting past the assumption that learning takes place in a course. Although we may have many mental models to update, a few examples illustrate the challenge.

Technology allows us to store massive amounts of information; the internet and search tools make it almost instantly accessible. If so, does memorisation lose some of its value? Are tomorrow's best learners those who can use the tools and technologies to sift, search and synthesise information or those who can memorise places, numbers, and formulas? Is knowing something or knowing how to make a decision more valuable? 'Will the ability to synthesise information become the primary goal of education?'⁵¹ In a world full of images, audio, and spatial relationships, do schools put less emphasis on reading and print and more on other skills?⁵² Is individual effort most important, or has thinking become a distributed activity – among people, devices, and digital resources?⁵³

If the most important skills in the 21st century are finding information and experts, engaging in complex communication, solving ill-defined problems, and making decisions, what are the implications for testing and assessment practices? Should exams be given without access to books, calculators, and the internet, or is the real measure of a learner's skill demonstrating how to use the tools and technologies to augment their own capabilities?

Other mental models also present challenges. Many of our students, and their parents, are focused on achievement: 'getting an A' so they can get a good job. If there is too much focus on getting the grade – and getting it as easily as possible – students may not be learning all they should. Problem-based learning methods and authentic learning models have been proven to be highly effective. However, students often complain that these alternative approaches require too much time. We need to help our students, parents, and communities see the value in these more complex learning environments. If instructors and students don't see the value of putting additional work into learning, it will be impossible to change the status quo or improve the competitiveness of the future workforce.



⁵¹ Futurelab (2007). *2020 and Beyond*. http://www.futurelab.org.uk/resources/documents/opening_education/2020_and_beyond.pdf Accessed September 19, 2007. Page 19.

⁵² *Ibid*, p.25.

⁵³ *Ibid*.

Our mental model for educational resources tends to be that we create something (such as a textbook) and it is fixed. What does developing academic material mean in a world of mashups and co-creation? Since the traditional tenure/merit and peer system is strongly linked to publication record, do you only step outside that model at your professional peril? Does it mean our students should be learning how to co-create or remix material rather than just authoring it? What advice does the academy have for Web 2.0 skills? The potential implications of Web 2.0 on education – and society – are extensive.

Conclusion

Our assumptions about students and what is best for their education may not be matched by today's reality. It is dangerous to assume that we understand students simply because we were once in the same shoes. Times change. Technologies change. Students change. And so does education.

Mobile, wireless, connected Information clouds and learning


Mark van 't Hooft Kent State University, Research Center for Educational Technology

God meant us to be wireless. The last cord we were connected to was cut at birth.

Frank Sanda, Japan Communications Ltd

Mobile devices and pervasive wireless networks are redefining the ways in which we live, play, work and learn, because they provide ubiquitous access to digital tools and 24/7 access to resources, tools, and communication channels in popular and widely-used portable forms. It is difficult for many of us to imagine life without the internet, and media sharing and social networking sites like YouTube, Flickr, MySpace, and Facebook have become household names in a relatively short period of time. When we're on the go we take our tools and access with us in the form of mobile phones, media players and other wireless devices, or we access technologies embedded in our physical environment.

We expect to have communication and information at our fingertips and customised to our needs. Think, for example, about tracking a UPS shipment online in real time, accessing restaurant reviews within physical proximity on your mobile phone, or communicating via SMS while commuting.



What is it that makes our current technologies so different? It is a combination of characteristics that:

- make the digital tools we use increasingly mobile *and* connected
- place control in the hands of the user and provide ample opportunities for personalisation
- enable us to have rich experiences as we access, aggregate, create, customise and share digital information in a variety of media formats, anywhere and anytime, and regardless of platform.

This article focuses on these three characteristics and the impact they are having on education and learning in an age in which the internet is not only always-on but also always-on-you. Mobile and wireless technologies provide opportunities for learning to become more personal and customised yet collaborative and networked, portable and situated, ubiquitous and lifelong. The rapid development and convergence of media, the read-write Web, and mobile tools and networks are opening up new opportunities for learning by allowing learners to be mobile, connected, and digitally equipped, no longer being tethered to a fixed location by network or power cables, a standardised curriculum, or a bell schedule. It also means that learning and formal education are increasingly at odds, as more and different types of learning are happening outside the classroom than in it. In sum, as our environment is becoming more flexible and unpredictable, so is our learning.¹

Mobility and connectivity

The proliferation of cheap mobile gadgets, wireless Internet access for everyone, a new Web built for sharing and self-expression... suddenly computing means connecting.

Wade Roush, Technology Review

In five years there will be one Internet and it will be mobile

Daniel Applequist, Vodafone

What does it mean to be mobile? Mobility can be defined as the ‘ability and willingness to move and change.’ [<http://en.wiktionary.org/wiki/mobility>] Mobility can be applied to many different aspects of our lives, such as social



¹ KnowledgeWorks Foundation (2006). Map of future forces affecting education: 2006-2016. <http://www.kwfdn.org/map/map.aspx>

or academic mobility, but for the purposes of this article I will refer to mobility as a person's physical movement from location to location, as is the case, for example, when commuting from home to work or school.

The proliferation of personal and mass transportation has made us more mobile than ever before, and developments in both hardware and software technologies are allowing us to remain connected to information and people in both the physical and digital realms. Some predictions indicate that by 2010, more than 3.2 billion people will own a mobile phone and subscribe to a cellular service.² A recent article by BBC News³ expects that the following five technologies will be making headlines in 2008: various forms of the web on the go (such as Google's Gears⁴ and Adobe's Air⁵ which provide offline functionality for web applications), ultra-mobile PCs (UMPCs such as the ASUS EEE⁶), IP-based TV, mobile VOIP (such as offered by Skype⁷, and Wimax (high-speed broadband over long distances). While the boundary between our physical and digital worlds is getting increasingly blurry, it does make sense to look at them separately first to better understand how our digital tools influence the ways in which we navigate each world and what implications this may have for learning inside and outside school.

Physical mobility and connectivity

It's about letting the computer get out of our way so we can work with other people and share our information.

Sam Schillace, quoted in Greene, K., 'Google's cloud looms large' in Technology Review 3/12/2007 [<http://www.technologyreview.com/Biztech/19785/>]

Smaller, faster, and wireless technologies coupled with constant improvement in wireless networks are allowing us to easily navigate our physical world *with* our tools. While the amount and variety of mobile computers, smart phones, digital cameras and media players are staggering and the possible combinations of functionalities into a single device even more mind-boggling, most people tend to use their devices for a limited number of tasks. The most important tasks include:



² Computer Industry Almanac (2006, March). Smartphones to outsell PDAs by 5:1 in 2006. Press release. <http://www.c-i-a.com/pr0306.htm>

³ <http://news.bbc.co.uk/2/hi/technology/7147804.stm>

⁴ <http://gears.google.com>

⁵ <http://labs.adobe.com/technologies/air>

⁶ <http://eeeepc.asus.com/global/product.htm>

⁷ <http://www.skype.com/mobile>

- accessing and aggregating information online (such as news, sports, entertainment, hobbies)
- navigating the physical environment (mostly by using GPS devices, Google Maps, and Google Earth for maps and directions, or sites like Flickr for geo-tagged images)
- accessing or interacting with digital information embedded in the physical environment (as available through (combining) technologies such as GPS, RFID or QR tags, and NFC (near field communication))
- interacting with the physical environment (for example, using embedded chips to make payments or accessing public transportation; a good example of this is the use of mobile phones for such activities in Singapore⁸)
- communication (using a variety of channels such as SMS, IM, VOIP, email and yes, even voice calls)
- entertainment (games, music, and videos/movies)
- media creation (recording video and audio, and to some extent, editing and publishing from a mobile device, for example by moblogging – uploading pictures from a mobile phone to a weblog)
- media tagging (labelling audio, video, and images with key words, or geo-tagging media for integration with tools such as Google Maps or Google Earth).

Digital mobility: cloud computing and the mobile web

New Web-based services don't just store your data online – they keep it synchronized across your laptop, desktop, and mobile phone.

Wade Roush, 'The Internet is your next hard drive', Technology Review, 24/7/2006

Desktop-based productivity and multimedia tools aren't the only game in town any more. Web-on-the-go applications are increasingly popular and becoming a viable and often free alternative, including tools such as Google Maps for Mobile, Flurry (mobile access to messaging, email, and feeds), and Xanco (one of many moblogs⁹). While many of these tools are web-based, others can be used both online and offline. Often, these applications allow you to work or



⁸ Bell, G., and Dourish, P. (2006). 'Yesterday's tomorrows: Notes on ubiquitous computing's dominant vision' in *Personal and Ubiquitous Computing*, 11(2), 133-143. <http://www.ics.uci.edu/~jpd/ubicomp/BellDourish-YesterdaysTomorrows.pdf>

⁹ <http://blogs.xanco.com>

access data offline when there is no connectivity, and when you get online they automatically sync with what you have stored there. Google Gears and Oracle 9iLite are good examples this type of hybrid use.

For example, the average Google employee keeps most of his or her software and data on the Web so that it is easily accessible from any location with internet access. This practice has become known as cloud computing, meaning that by using a device as simple as a UMPC or mobile phone and the internet, one can access a 'cloud' of computers holding massive amounts of information and offering large caches of data storage space, as opposed to a limited personal computer or application server.¹⁰ It also means that many of us are now creating, editing, aggregating, mixing, storing, and sharing 'clouds' of information on the internet. Our data and services are increasingly device-independent and platform agnostic, and accessible no matter where we are or what device we use. As long as we can connect to it, the internet will follow us. Thus, we keep our calendars online, and create, edit, and collaborate on written documents, spreadsheets and diagrams using Google Docs and Gliffy. We edit our pictures with Picnik before posting them on Flickr and share our movies on YouTube. We create and share custom maps using Google Maps and Google Earth. We tag our creations in a variety of ways so that others can find them, and we create personalised portals to digital information using RSS aggregators, and customisable browsers like Flock and mob5. Finally, we can synchronise all of this information across multiple devices using services such as Streamload's MediaMax.

In sum, cloud computing is fundamentally changing the way in which we use the internet in our daily lives. Thomas Vander Wal, Principal and Senior Consultant for Infocloud Solutions, describes the shift in use as going from the "I go get web" (people accessing static and information in proprietary formats created by others on a desktop computer) to the "come to me web" (people creating, finding, using/re-using, sharing and storing information in open formats across multiple devices in different locations). Consequently, whereas the focus used to be on the technology, it has now shifted to the person, demoting the technology to a serving role and following the user wherever he or she goes.¹¹

Also, information clouds come in different types and sizes, and information can be moved from one cloud to another (with the exception of the external infocloud). While I have only described the personal infocloud so far, Vander Wal has defined four different categories:¹²



¹⁰ Greene, K., 'Google's cloud looms large', *Technology Review*, 3/12/2007.
<http://www.technologyreview.com/Biztech/19785>

¹¹ http://personalinfocloud.com/2006/01/the_come_to_me_.html

¹² <http://www.vanderwal.net/essays/moa/040608/040608.pdf>



Four types of information clouds. Graphic courtesy of Thomas Vander Wal, <http://personalinfocloud.com> and <http://www.infocloudsolutions.com>

- **Personal:** a user's aggregated information repository that is controlled and organised by the user. It is portable and therefore easily accessible and re-accessible by its creator. Examples include personalised portals, PIMs (Personal Information Managers), data repositories (such as a calendar), and online wishlists.
- **Local:** a location or membership-based network, only accessible by a select few on a LAN or intranet. Not user-controlled or organised, but the structure is often familiar to the user. Examples of local infoclouds include social software, affiliations (work, organisations), and friends.
- **Global:** a.k.a. the internet, which can be accessed from anywhere and is not user-organised or controlled.
- **External:** information repository (location-based) a user does not have access to because it is behind a firewall and the user is not a member, resulting in limitations on information sharing.

One important characteristic of cloud computing is the mobility of the user and his/her tools, accessing and manipulating the information cloud whenever and wherever it is necessary or convenient. This is why discussions about the mobile web and what it should be have become much more heated as of late. There are those who see the mobile web as no more than pushing the internet onto mobile devices. Their main concern is figuring out ways to get existing internet content to fit on a smaller screen. A substantial part of the W3C Mobile Web Best Practices 1.0 document deals with this issue.¹³

Others are pointing to the importance of open standards and platforms. A major development in this area is Google's Open Handset Alliance and Android, an open and comprehensive platform for mobile devices. Moreover, Google announced that their goals in this area will be independent of device or platform, opening the door for many new and innovative applications that will take full advantage of the mobile devices they will run on.

Yet others still say the mobile web is much more. For example, according to Russ McGuire, director of corporate strategy for Sprint, the always-on nature of mobility, mobility as the first truly personal information technology with mass adoption, and mobile devices as having greater visibility into the context in which they are being used than previous technologies are driving this development.¹⁴ It is the *combination* of mobility and access that is important to this group of people, a combination that is important to education as well, as it provides opportunities for situated learning in ways not possible before.



¹³ <http://www.w3.org/TR/mobile-bp>

¹⁴ <http://law-of-mobility.com/2007/10/17/mobile-20-mobile-internet-vs-mobility>

Finally, any discussion of mobility, connectivity, digital information and education would be incomplete without mentioning the importance of context. Context provides meaning to information, whether it is a physical or digital location, a task, other information, a collaboration, work vs personal life, or something else. While we tend to use different types and pieces of information in different contexts, when we use the same piece of information across contexts it can have different meanings. Think for example how the meaning and use of historical information changes from reading it in a classroom to accessing it at its historical location. Second, context is never static, and shapes (and is shaped by) digital technologies, especially those that combine physical locations with digital information. Third, context enables users to establish at least a sense of control over their digital tools and information (and vice versa) by the choices they make in what they access and when they access it.

As a result, context/location-aware technologies are becoming a major trend. GPS devices are a prime example of this development, as they recognise where you are and can deliver relevant information based on your location, such as directions to a destination. However, there are many more technologies that use location as an anchor point for digital information, including:

- **object recognition:** users can take a picture of a building or object in their environment with their mobile phone, and receive information about that site, potentially turning every digitally tagged object in one's surroundings into a learning opportunity. Mapion Search in Japan does exactly that, while Nokia has been talking about the development of similar Point and Find technology. Another example is Nokia's Shoot to Translate, which will allow users to take a picture of text in a foreign language, and receive a translation in real time.¹⁵
- **user recognition:** researchers at Xerox PARC are developing a mobile software application called Magitti, which makes recommendations to its users based on a variety of cues, including time of day, user behaviours, and text messages.¹⁶
- **geo tagging:** this has become a common practice on sites like Flickr, and allows users to not only place pictures within their geographic context, but also have search results displayed in that way. For educational purposes this creates all kinds of interesting opportunities, such as mapping existing habitats of certain mammals or birds, or virtually visiting a location by searching for and organising geotagged images of that location



¹⁵ http://www.nokia.com/NOKIA_COM_1/Press/twwln/press_kit/TWWLN_umbrella_release.pdf

¹⁶ Greene, K., 'New software uses artificial intelligence to infer your behaviour and serve up appropriate lists of restaurants, stores, and events', *Technology Review* 110. <http://www.technologyreview.com/Biztech/19698>
To see how Magitti works, see <http://www.technologyreview.com/player/07/11/13Greene/1.aspx>

User control and personalisation

The real Web 2.0 is about control. It's about letting each user control their own interactions, and that is why people like it. At the end of the day, Web 2.0 is about being selfish, and the projects that will succeed are the ones that are embracing that fact.

Rob May, University of Louisville [http://www.businesspundit.com/50226711/punishment_selfinterest_and_cooperation_why_the_altruism_of_web20_doesnt_exist.php]

The personal Internet of tomorrow will serve you – delivering the information you want, when you want it, how you want, wherever you are.

Paul Otellini, President and CEO, Intel Corporation [<http://www.intel.com/pressroom/archive/releases/20080107>]

A second important characteristic of current digital technologies is that they put an increasing amount of control in the users' hands. What this means is that users have more (and often free) tools and opportunities to build, rebuild, and modify personal infoclouds. It also means that users have more of a say in when, where, and how they access information and communicate with others. This is especially the case when Web 2.0 technologies are being made accessible on mobile devices, enabling users to manipulate and customise the various information clouds on the fly. A good example is YouTube's recent announcement that it will widen its mobile service to a broader range of devices.¹⁷

User control is becoming increasingly important. A recent study by Nokia and The Future Laboratory¹⁸ concluded that by 2012, up to 25 per cent of 'entertainment consumed by people will have been created, edited, and shared within their peer circle rather than coming out of traditional media groups'. Associated opportunities for learning are substantial and can put learners and educational institutions at odds as they compete for control over learning.

However, it also means that once digital information is created, whether it be a web page, a video, or a series of images, this information can be re-used, changed, mixed, re-organised and shared by others. While this creates opportunities for those who find and use digital content, it also forces the original creator to relinquish at least some of his/her control. In addition, even though users benefit from sites that are built through user contributions, some of the benefit goes to the site itself (and its advertisers).



¹⁷ http://www.youtube.com/blog?entry=sAyN42I_HMI

¹⁸ <http://www.thefuturelaboratory.com>

In sum, increased control comes with a cost. I will return to this issue later on, as it has implications in the realms of copyright and privacy. In the meantime, it is important to remember that in many ways, 'mobile 2.0 is not 'the Future'. It is services that already exist all around us. These services are maturing at an amazing rate and what they are doing is effectively knitting together Web 2.0 with the mobile platform to create something new: a new class of services that increase mobility but are as easy to use and ubiquitous as the Web is today. These services point the way forward for the mobile data industry.¹⁹

Rich experiences

Web 2.0 is based on user intelligence instead of technologies, i.e. by giving users smart tools that enable them to apply human semantics to information provided, you get a more intelligent web.

Enrique C. Ortiz, CTO, eZee Inc., as quoted in http://www.readwriteweb.com/archives/understanding_mobile_2.php

Web 2.0 has been defined by Tim O'Reilly as:

[the]network as platform, spanning all connected devices; Web 2.0 applications are those that make the most of the intrinsic advantages of that platform: delivering software as a continually-updated service that gets better the more people use it, consuming and remixing data from multiple sources, including individual users, while providing their own data and services in a form that allows remixing by others, creating network effects through an 'architecture of participation,' and going beyond the page metaphor of Web 1.0 to deliver rich user experiences.²⁰

In other words, developments in web applications in recent years have allowed users to do more than just consume information in a fixed location. We can now:

- access information within a local context (augmenting reality with digital layers, tagged and hyperlinked environments, QR codes, and RFID tags that can be captured with a mobile camera phone so that users can retrieve additional information about the object that's been tagged.²¹ Examples of this range from Japanese beef sporting QR tags that provide customers with a



¹⁹ From Daniel Applequist's post "What is mobile 2.0 (beta)?". <http://www.torgo.com/blog/2006/11/what-is-mobile-20-beta.html>

²⁰ http://radar.oreilly.com/archives/2005/10/web_20_compact_definition.html

²¹ See for example Jonietz, E., 'TR 10: Augmented reality', *Technology Review* 110(1). <http://www.technologyreview.com/Infotech/18291/?a=f>

complete history of the product to tagged urban environments transformed into spaces for public play.)²²

- aggregate information from multiple sources (by way of RSS, aggregators, customisable browsers, and portals)
- create new information (using blogs, wikis, Jumpcut, Fauxto, Picnik, etc.)
- collaborate with others (Google tools, Gliffy, wikis, folksonomy, etc.)
- customise content to our likes and needs (by social bookmarking, folksonomy, mashups, using portals, etc.)
- share what we've created (using YouTube; Flickr and so on)
- network and communicate online (by way of MySpace, Facebook, IM, SMS, VOIP, etc.)

Learning while mobile and connected

Mobile learning is the processes of coming to know through conversations across multiple contexts amongst people and personal interactive technologies.

Mike Sharples, Josie Taylor, Giasemi Vavoula [http://telearn.noe-kaleidoscope.org/open-archive/browse?resource=215_v1&back=%2Fstats%2Flast_week_popular_publications.php]

How do the three characteristics of mobility/connectivity, user control, and rich experiences affect learning? The combination of mobility and connectivity provides new opportunities for learning at times and in locations not previously considered practical or useful. It has also created increased possibilities for situated learning. User control has often put learners and formal educational institutions at odds, especially when learners provide their own devices. As schools and learners vie for control of the learning process, a common result is for students to do most of their real learning outside the classroom, as schools limit access to many of the digital tools that students are accustomed to using in their daily lives. This trend has also empowered learners to construct customised PLEs around topics of interest, enhanced by rich media experiences as they access, aggregate, mash, manipulate, create and share information from a wide variety of sources near and far.

To illustrate this, let's take a look at a couple of learning scenarios that incorporate mobile and connected learning. They range from formal to informal.



²² See for example <http://www.playareacode.com/index.html> and <http://homepages.nyu.edu/~dc788/conqwest/about.html>



Sample screenshots from MyArtSpace mobile phone interface. Courtesy of Giasemi Vavoula, Mike Sharples, Paul Rudman, Peter Lonsdale, and Julia Meek.

MyArtSpace (United Kingdom)

MyArtSpace²³ is a service for children to spread their learning between schools and museums. It currently runs in three UK museums: the Urbis Museum of Urban Life in Manchester, the D-day Museum in Portsmouth, and the Study Gallery in Poole. It can be used for informal learning, but is best suited for school field trips. The aim is to make a day out at the museum part of a sequence that includes discussing an open-ended question in the classroom, exploring it through a museum visit, reflecting on the visit back in the classroom or at home, and presenting the results.

The technology used (mobile phones and personal web space) provides the essential link across the different settings. Learners use mobile phones to collect multimedia representations of exhibits and store them online; to collaborate with those who have collected the same digital artefacts as prompted by those same mobile phones; and to reflect on what they see, hear, and discuss. Students can access their personal collections online after the museum visit, modify them, and create web-based galleries to share with others at school and elsewhere.

This scenario illustrates how learning in a formal setting (school and museum) can be amplified and extended when learners generate and manipulate content within a meaningful context, communicate with others and their environment using mobile devices, and share what they learn with others. In addition, the MyArtSpace project also shows how digital content can be device independent and be accessed whenever and wherever the need arises. In contrast, there is definitely room for improvement, as learners are limited in the physical and digital content they can access (that is, what is provided by the museum), and how they access it at particular times (mobile phones are provided by and only used in the museum). While these limitations do help provide a certain amount of focus, it should not be too difficult to extend the project even further by incorporating the use of (mobile) Web 2.0 tools to create, aggregate, manipulate, mix, and share additional digital content. For example, students visiting the D-Day museum could also visit nearby physical locations of importance or interview D-Day veterans or family members who experienced the event.

National Museum of Natural Science (Taiwan)

An example of more informal and personalised learning while mobile is the context-aware guiding service in the National Museum of Natural Science (NMNS) in Taiwan.²⁴ Before going to the museum, visitors can log onto its website, create an itinerary that fits individual needs and interests, and save



²³ <http://www.myartspace.org.uk>

²⁴ http://www.nmns.edu.tw/index_eng.html

these preferences in the museum's database. At the NMNS, visitors are then equipped with internet-ready wireless handheld devices, giving them three options: following the individually prepared plan, joining a recommended learning tour, or freely exploring exhibits. The context-aware system automatically determines the visitors' location and delivers corresponding content and relevant information to their mobile devices. After the visit, the web-based system provides additional learning content and recommends further resources according to the individual visitor's on-site learning behaviour and preferences.

Even though the museum experience here may seem similar to the one in the first scenario with regard to using a mobile and networked device to access location-based, digital content, it is very different in that it is the learner who decides what to learn about, and whether to do this by following a recommended tour, a personal plan, or by unstructured exploration. However, just like in the first scenario, the experience would be more powerful if learners could bring in their own mobile wireless devices that they are familiar with and that are customised to their preferences and uses.

In addition, it would be easy for a museum visitor to expand his/her learning following the visit by looking at additional learning content recommended by the museum's web-based system and combining it with content accessed and aggregated online. The user would decide what direction to take, what to look at, create, share, etc., and how much time to spend on this learning expedition, without having to worry about a teacher or a set curriculum. By creating a context that combines both physical and digital spaces, learners can construct what has been described as a Personal Learning Environment (PLE) which allows them to consume as well as create learning resources.

PLEs can be applied to many different educational contexts. For one, students in formal learning environments could easily create a PLE for a research project, accessing, aggregating, organising, mashing and manipulating available resources, and sharing what they've learned. The mobility and online accessibility of PLEs is also very useful for learning contexts in which learners are not confined to a particular learning space or are physically removed from their teachers/mentors. For example, in New Zealand, trainee bakers are using a combination of mobile and web-based tools including mobile phones, Flickr, YouTube, Springdoo and Moodle to access content and assessment, and create e-portfolios. The e-portfolios are used to track learning, create working portfolios, get feedback from peers, university mentors, and employers, and provide authoritative evidence of a trainee's knowledge and skills.²⁵ Another example of a PLE that supports on-the-job training is Mobiletools' eTaitava, a



²⁵ Chan, S., and Ford, N. (2007). *mLearning and the workplace learner, Integrating mlearning eportfolios with Moodle*. Paper presented at MoLTA 2007. <http://molta.massey.ac.nz/massey/fms//Molta/Chan.pdf>

multimedia reflection and feedback tool on mobile phones that connects learners, teachers and workplace mentors in on-the-job training environments. Its strength lies in its capability to provide all parties involved with continuous and immediate feedback on learning progress of trainees.²⁶

Writing this article

Believe it or not, writing this article is probably the best of the three examples of what learning while mobile and connected could look like. While the example I describe here describes an office-type situation, just like in the previous scenario it wouldn't be too difficult to apply it to student learning in formal learning environments. In this scenario we would replace the manuscript with a student project, and the author of the manuscript with a student accessing various resources while being in multiple locations and using a variety of digital tools.

As I set out to write the manuscript, I worked on an outline with the editor from Becta, who is literally thousands of kilometres away and whom I've never met face to face. I wrote the article in various locations that are more than 2,500 kilometres apart physically, and wrote during the day as well as in the middle of the night, storing materials both locally and online. I accessed resources on the internet as I was writing, scanning a substantial number of sources using Feedreader, an RSS feed aggregator, and conducting many Google searches that were both intentional and unintentional, as new questions arose when others were answered. I also emailed scholars and industry experts in two countries in order to collect images for use in this article (a couple of which I did some quick editing on using Picnik). Finally, I IM'd with one of my colleagues about the manuscript and worked with several editors at Becta on the various drafts to create the one that you are reading now. Most likely, you are reading it either directly on the internet or as a copy downloaded on your desktop or mobile device, in a format that is open (PDF) and using software that is free (Adobe Reader).

Again, this example shows how easy it is to set up a digital and customised personal learning environment that will follow you wherever you go, as long as you have internet access. If I had really wanted to, I could have worked completely online, using web-based applications (such as Writely in Google Docs) and stored files online in my personal infocloud. While my learning experience was not necessarily situated, it was definitely personal and customised, collaborative and networked, and ubiquitous. I did not have to worry about not having access to the resources I needed to write the article, even when being thousands of kilometres from home.

Issues and barriers

Learning while mobile and connected and storing your digital information online does not come without its drawbacks. As difficult as it is going to be to implement more fully a concept of learning that moves control from teacher to learner, there are some additional issues that need to be addressed. In many cases, they've already reared their heads. These are the most important challenges:

- Privacy and online safety: when storing your personal profile and information online you give up something to get something. Because many popular web-based services are free or almost free, and we expect them to be, we end up trading in a piece of our privacy in exchange for a service, that is, privacy becomes the currency. Google plays on this by tracking our moves on the web and tailoring its advertising to our online behaviour. Facebook's Beacon is an example of an even more invasive advertising system as it collects 'information about user actions on affiliate sites regardless of whether or not the user chose to opt out, and regardless of whether or not the user is logged into Facebook at that time.'²⁷
- A related issue, that of authenticity. Many of us do not provide complete or true information when creating a profile on a social networking or media sharing site, and it is easy to hide behind a fictional username and picture. We can only guess how many kids have MySpace pages who are under the age limit set by the site, as it is pretty much impossible to verify the age of each and every MySpace user. In addition, many profiles on social networking sites are spoofs, as plenty of school administrators can attest to and often leading to nasty situations.²⁸
- Data security: even if we can post our pictures, movies and profiles online without giving up our privacy, how do we know that the service we use is secure? Again, it all comes down to issues of control, which leads us to
- Ownership of data: who owns the information? The creator? The service that stores it? The user who finds, aggregates, and remixes it? All of them? Some of them? None of them? In an era in which digital information can be effortlessly created, duplicated, changed and shared, ownership of content is becoming increasingly tricky. Traditional copyright laws are difficult to uphold in an arena that crosses national boundaries and jurisdictions, and where the concept of copyright itself may mean different things in different countries. Instead, it seems as if more and more creators of digital content



²⁷ <http://community.ca.com/blogs/securityadvisor/archive/2007/11/29/facebook-s-misrepresentation-of-beacon-s-threat-to-privacy-tracking-users-who-opt-out-or-are-not-logged-in.aspx>

²⁸ See for example the saga of a student in North Carolina who was suspended for ten days for posting an altered picture of his school's assistant principal on MySpace. http://www.splc.org/report_detail.asp?id=1256&edition=38

on media-sharing sites are moving to a Creative Commons copyright, which has different levels, but basically states that others may use, redistribute, and in many cases modify your work, with whatever restrictions you choose to put on it.

- **Connectivity:** the concept of infoclouds only works if ample broadband connections are available and wireless coverage is continuous and consistent. Even today, attaining uninterrupted wireless connectivity over a large geographic area is still a costly proposition, although alternatives such as mesh networks (which are possible, for example with the OLPC's XO device), or learners' personal connectivity are being considered.
- **Ownership of devices:** currently, most educational initiatives that use mobile technologies provide the devices to learners (see for example the MyArtSpace scenario above). Evidence shows that in the long run this is an untenable situation, as many schools simply cannot afford to provide each learner with a mobile device. Even when devices are donated to schools through grant funding or donations, projects often die when the time comes to replace worn out and obsolete technologies. An alternative that is being considered by some is using the mobile technologies that learners are already carrying in their pockets. Can education make use of these devices by merely providing wireless network access in learning spaces such as schools or museums? And if this works, how do issues of privacy, online safety and data security change? What shifts in responsibilities and liabilities will we see?²⁹

In addition, issues that are more directly related to learning include the following:

- **Relevance of technology:** are we using digital tools as a means or an end? Are we using technology because it amplifies learning in ways that we can't do without or are we just following the latest fad?
- **Issues of technology and media literacy:** just because learners know how to use digital technologies for personal use or entertainment doesn't mean that they'll automatically know how to use these same tools well for learning. We need to look more closely at the cognitive and literacy demands of common media tools that are placed on learners when asked to perform a good search, interpret text and graphics, or identify and judge a source of information.³⁰



²⁹ A good example that shows the complications when student-owned devices are involved is the cell phone porn scandal at a Pennsylvania school, in which pornographic video and photos of two US high school girls were transmitted by cell phone to dozens of the girls' classmates and then to the wider world. According to one student, "The school isn't going to get everybody because it is everybody. I don't know anybody who didn't get the pictures," said a 16-year-old junior who said she deleted the images when she got them. For the full story, see <http://www.msnbc.msn.com/id/22840727>.

³⁰ Tally, B. (2007). Digital technology and the end of social studies education. *Theory and Research in Social Education*, 35(2), 305-321.



- A second issue of media literacy has to do with the safe, ethical, and responsible use of the internet. Although students can be very technologically savvy, kids of all ages need to learn how to behave appropriately and protect themselves on the internet. Since virtually any post on the internet leaves behind a digital trace, it is open to almost anyone with the means to find it in cyberspace. What teenagers often do not realise is that what they post or upload for friends may be read by school administrators, college admissions officers, potential employers, or even internet predators.
- Assessment of learning: how does the use of personal information clouds for learning change who assesses what, when, where, and how?

Conclusion

It should be obvious by now that the use of mobile and wireless technologies in combination with personal information clouds that promote learning that is personal and customised yet collaborative and networked, portable and situated, ubiquitous and lifelong, is on a collision course with those formal institutions of learning that are trying to hold on to a model of learning more fit for 1908 than 2008. Mobile, wireless, and connected technologies are here to stay. In fact, new technologies that are always with us, allow us to create more detailed and true-to-life online identities, understand our preferences, and adapt more closely to the chronology of our lives are making computing an integrated part of our lives. Wade Roush has called this continuous computing, and likens the experience to 'wearing eyeglasses: the rims are always visible, but the wearer forgets she has them on – even though they're the only things making the world clear.'³¹

The key for education is that these technologies are providing opportunities for learning that don't really exist without them. Instead of banning these tools and punishing those who try to use them for learning in formal settings, schools should take a serious look at learning more about the digital technologies so prevalent in many areas of our society today.

Should schools be afraid that learning while mobile, using the mobile internet and infoclouds, is going to replace formal approaches to learning in the near future? Probably not. However, institutions of education should be ready to at least consider using mobile, wireless and connected tools as supplements to what's happening in the classroom, in an effort to make learning both more personal and connected to the real world. Wireless mobile devices could



³¹ Roush, W., 'Social machines' in *Technology Review* 108(7) 44-49. <http://www.technologyreview.com/Infotech/14664>

provide the bridge to make that happen, as long as schools provide the networks and the access, because in the long run, the learners will provide the devices to access their applications and information stored online.

As Chris Lehman, founding principal of the Science Leadership Academy in Philadelphia has said, 'Our world is changing, and the changes don't make school easier. If anything, it makes it harder, because we can't pretend there's a clear-cut roadmap ... Our schools must be personal, they must be community-based – however we choose to define our community. They must be relevant, and they must be willing to change. Our students, the citizens of this new century, deserve nothing less.'³²



³² Lehman, C. (2007). School 2.0. The Science Leadership Academy. *Learning and Leading with Technology*, 34 (7). Available at http://www.iste.org/Content/NavigationMenu/Publications/LL/LLIssues/Volume_34_2006_2007/April_No_7_2/LandL_April_2007.htm, which for many of you will be part of an external and therefore inaccessible information cloud...


Location-based and context-aware education: prospects and perils

Adam Greenfield [New York University](#)

The roots of ubiquity

For so very many of us, the personal computer has become one of the few truly indispensable tools of life. In little over a decade, the networked PC has subsumed into its universe of ones and zeros the signature tasks of an increasingly wide panoply of everyday artefacts: the stereo systems and day planners, alarm clocks and typewriters with which we formerly furnished our lives. We use PCs for everything from scheduling our personal and professional commitments, to keeping in touch with far-flung friends and family, to watching movies and playing games.

But we know that technology, and particularly information technology, is never static. What might come after the PC? And what consequences might arise as a result of the shift away from these devices (and the ways in which we use them) and toward some as-yet unknown combination of technology and practice?



In the early 1990s, the late Mark Weiser, then Chief Technologist at Xerox's legendary Palo Alto Research Center, began articulating the first coherent vision of information technology in the post-PC era, a vision that would become known as 'ubiquitous computing'. Taking the continuing validity of the engineering and business axiom known as Moore's Law – in which microprocessors' power doubles every 18 months to two years, while size and cost are halved – as a given, Weiser described a world where such processors would be embedded in many of the ordinary objects and surfaces of everyday experience. Computers *per se* might disappear, but their ability to sense, manage, represent and act on information would become immanent, 'invisible, but in the woodwork everywhere around us'.

Fast-forward fifteen years, and it is clear that while Weiser may not have been correct in all respects – particularly, his depiction of ubiquitous computing as primarily transpiring in purpose-built, heavily-instrumented spaces has been subsumed by our casual embrace of mobile phones and wireless internet – the broader contours of his vision are indeed coming to pass. Information processing is increasingly escaping the confines of the desktop and permeating the everyday world, showing up in the widest possible variety of consumer products: from Oyster cards to in-car satellite navigation systems to Nike+ bio-telemetric running shoes, the possibilities of the everyday are being decisively transformed by this encounter.

Toward the locational and the contextual

One of the consequences of Moore's so far unstoppable march forward has been that GPS and other location sensors are cheaper, smaller, and more robust with every passing quarter, continually easier to install and to use. Thanks to the ability to deploy such devices in multi-sensor packages capable of cross-checking satellite positioning against triangulated Wi-Fi base-station or cell tower strength, they are also increasingly accurate. Simply put, we are able to endow an ever-greater proportion of the things we devise with the ability to locate themselves in space and time, to whatever degree of precision is required.

At the same time, we see increasing use of relational 'inference engines' able to model circumstances at a higher level of abstraction than that provided by raw sensor input, by applying heuristics to that input. For example, a ubiquitous building-management system might infer from a certain fact pattern – say, that I am currently seated in my task chair, my word-processing application is active, and my office door is closed – that I do not wish to be disturbed, and accordingly route incoming calls directly to voice mail.

When taken together, these two capabilities underwrite the provision of a sophisticated class of informational services, variously thought of as

‘location-based’ or ‘context-aware’. In such applications, the value proposition to the user hinges on the provision of ‘the right information, at the right time, in the right place’.

This could mean something as simple as a taxi’s cross-referencing its current coordinates with a networked database of bank branches, in order to display the location of the nearest ATM on its rooftop advertisement. It could be the hackneyed ‘m-commerce’ scenario, all but invariably trotted out in these discussions, of a discount coupon sent to your phone whenever you pass through the ‘catchment area’ of a Starbucks or a McDonald’s. Or it might simply mean that the information pushed to you varies with where you are, who you happen to be with, and what tasks you’re currently engaged in.

At least in theory, this means an experience of effortless utility, with users able to access networked information without having to subject themselves to the various hassles and compromises that attend the general-purpose PC. The examples given by researchers in the field are generally self-consciously drawn from mundanity, from the shared circumstances of everyday life: the day’s weather displayed on your bathroom mirror on awakening; the traffic report overlaid onto your windshield; the subtle cue embedded in your wallet or handbag that lets you know when one of your friends is within a hundred metres of your present position, and available for socialising.

It hardly needs to be said, though, that one of the domains of practice most often mentioned in connection with both location-based and context-aware services has been education. The idea that ubiquitous technology might circumvent some of the institutional failings of education as it is presently generally constituted in the developed world, allowing learners to build knowledge at their own pace, on their own terms, and with a high degree of vividness, is remarkably seductive. Realising such promises in any meaningful sense, however, may take a good deal more effort than has previously been acknowledged by the more techn-utopian proponents of ubiquitous computing.

The complications of ‘context’

Despite their common technical underpinnings, there are important distinctions between location-based systems and putatively context-aware ones, and these in turn have significant implications for the suitability and short-term deployability of such systems as educational interventions.

The fundamental premise of so-called ‘location-based learning’ is that, as geographer and spatial-information specialist Anja Kipfer expressed it in a 2006 article, ‘building a spatial and thereby visual or haptic connection between a learning object and the learning content stands for a better cognition and remembrance’ [sic]. (The location-based rubric subsumes mobile interfaces

to local content, as well as a class of systems that architects are increasingly beginning to think of as ‘situated technologies’, that is, installations or interfaces otherwise permanently embedded in place.)

However technically difficult this might be to achieve and to implement with any subtlety, it’s at least a relatively straightforward proposition: place plus narrative equals a more resonant and memorable experience than either place or narrative alone. So far, so good – and hard to argue with. But as we shall see, true ‘context awareness’ is an entirely thornier prospect.

The notion that ubiquitous technical systems will ever be endowed with sufficient discernment to parse and comprehend events transpiring in their environment is far from uncontroversial. As ubiquitous-computing researchers Tim Kindberg and Armando Fox point out in a 2002 paper, ‘[s]ome problems routinely put forward [as suitable for context-aware systems] are actually AI-hard.’ That is, these are challenges of the same order of complexity and difficulty to the successful creation of an artificial human-level intelligence. Irrespective of the advance of Moore’s Law, none of the technical development that has happened over the five years since has done anything to alter this fundamental condition.

Such intractability, it turns out, is anything but a rare property in the daily and the ordinary. There are situations aplenty in everyday life the decoding of which we rightly regard as trivial, but which remain beyond the ken of either brute-force computation or more advanced neural network pattern-recognition systems. The example of intractability Kindberg and Fox offer – whether a technical system can accurately determine if a meeting is in session in a given conference room, based on the available indicators – is startling at first blush: a split-second peek through a cracked door is generally sufficient for us to verify whether or not what’s going on in the room constitutes a meeting. How could something so manifestly obvious to us be so difficult for even the most sophisticated digital systems to make sense of?

As best we understand it, that split-second determination of *gestalt* is something we build up from dozens, even hundreds of subtle, implicit cues. We start by noting the same fundamentals that a grid of networked ubiquitous sensors might pick up: is there anybody in the room at all? Is there more than one person present? Is anyone speaking? But right away, we begin to depart from the sort of evidence that can be adduced and quantified by even the subtlest of sensors, at least as yet. We register the body language of participants, the pitch and volume of their voices, their choice of topics and words, even factors like the presence or absence of tools for note-taking, and from these observations we build our judgement. We build it exceedingly swiftly and, happily, we are rarely wrong.



The conference-room example can be supplemented with many others. Knowing when a loved one's feelings have been hurt, when a baby is hungry, when confrontation may prove a better strategy than conciliation: these are things that human beings of even moderate sensitivity are generally able to detect in short order, but even the most sophisticated contemporary inference engine cannot determine with any consistency at all. It turns out that when it comes to human interactions, 'context-aware' systems are, frankly, best described as autistic.

Values here

Nor is this the only issue that confronts the would-be employer of ubiquitous systems in the educational milieu. There is also the problem that the majority of such projects have so far been proposed and devised by engineers – by human beings, in other words, with a very high degree of technical competence, who are generally not nearly as grounded in the social sciences, and may indeed regard same with suspicion, disdain or outright hostility. (This is purely a matter of my own observation and opinion, but I would even venture that many engineers immersed themselves in the technical domain in the first place because the certainties they found there were more comfortable for them than the messy interpersonal world.)

The result can be, and all too frequently is, a horribly clunky reification in the systems they devise of models of human cognition and learning that very few people in the social sciences would recognise as current or complete. Particularly, there has been a historical tendency in ubiquitous-system design to map strategies for educating people on the ways in which technical systems are trained.

Illustrating the point vividly is the abstract of a recent paper entitled 'A Ubiquitous Education Support Method based on Analysis of Learning Patterns Using Rule-Based Reasoning' (quoted verbatim):

Many support methods based on learner's learning behaviours and psychological states have been proposed with the progress on Web-based learning. However, the learning in real world is the most important part in one's learning. Researches on learning support methods in real world become possible and necessary with the progress on the ubiquitous techniques. In this paper, we propose a hierarchy of zones, by which the position information and learning actions of learners can be correctly caught, and the services can be easily provided to learners. And a method to analyze learners' learning patterns from learning histories based on learning orders and reactions in hierarchical zones, and provide supports to learners by using Rule-Based Reasoning (RBR) in order to increase learning efficiency and help learners to bring up good learning (life) rhythms.

This is clearly how an inference engine works, but among the developmental psychologists I've spoken to, at least, the jury is still out as to whether or not it's how human beings form judgements about the world. The desire to provide appropriate cues to different kinds of learners is commendable, but in this case the specific architecture chosen to address the task seems overly mechanistic.

Recalling the example of the conference room, it is highly likely that learners signal their true readiness and capacity to absorb new knowledge with a wider and more subtle variety of behaviours than are accounted for in the clearly delineated, yes/no 'hierarchical zone' protocol proposed here. The danger is that because such models are easier for engineers to reckon with – easier, that is, to translate into the crisp, binary certainties required by digital information-processing systems – they will have more appeal to developers than the detailed, painstaking, difficult-to-generalise ethnographic work that might support better tools.

The truly problematic aspect of the paper, though, is contained in the abstract's last sentence, for what it turns out is being proposed here is nothing less than a system that infuses its users with 'good life rhythms', without anywhere specifying who believes those life rhythms to be desirable or how they arrived at that determination. It is when such valuations get encoded in educational systems beneath the threshold of either their developers' or their users' conscious awareness that the most vexing issues arise. And this is most especially the case when such systems are deployed, as experience has taught us will often be the case, in populations other than those for whom they were originally designed.

Take, for example, JAPELAS, a prototype context-aware ubiquitous Japanese-language instruction system developed at Tokushima University (2005). One of the complications of learning to speak Japanese involves knowing which of the many levels of politeness is appropriate in a given context, and this is just what JAPELAS sets out to teach. JAPELAS determines the apposite expression for a given situation by trying to assess both the social distance between interlocutors, their respective status, and the overall context in which they're embedded. In this model of the world, context is handled straightforwardly as being a matter of where you are physically: is the setting a bar after class, a job interview, or a graduation ceremony? Social distance is also depicted as being relatively simple to determine: are my interlocutors students in my class, in another in the same school, or do they attend a different school altogether? But in order to gauge social status, JAPELAS assigns a rank to every person in the room, and this ordering is a function of a student's age, position, and affiliations.

Observationally, none of this appears to raise any red flags for the Japanese. There is, for that matter, little reason that it should: all that JAPELAS does is encode into a particular technical system rules for linguistic expression that are ultimately derived from conventions about social rank that already existed


in the culture. Any native speaker of Japanese makes determinations like these a hundred times a day, without ever once thinking about them: a senior 'outranks' a freshman, a teaching assistant outranks a student, a tenured professor outranks an adjunct instructor, and a professor at one of the great national universities outranks somebody who teaches at a smaller regional school. To Japanese people, such orderings are natural and obvious.

But for more than a few non-Japanese considering such systems, including myself, the inscription of distinctions like these in the unrelenting logic of an information-processing system is a source of substantial unease and discomfort. Admittedly, JAPELAS is 'just' a teaching tool, and a prototype at that, so maybe it can be forgiven a certain lack of nuance; would-be speakers of Japanese can expect to be drilled with much the same set of rules by just about any human teacher (I certainly was). It is nevertheless disconcerting to think how easily such distinctions can be hard-coded into something seemingly neutral and unimpeachable, and to consider the force they have when uttered by such a source.

This force is redoubled by the very pervasiveness of this class of applications. Where PC-based learning systems of the past also observed similarly problematic distinctions, they at least did so in their own bounded frame of reference. The same property of ubiquitous applications that endows them with unprecedented vividness and impact – their taking effect 'out here in the world' with us, as opposed to being delimited to the confines of the screen – can lend the social and other orderings latent in them an almost untoward intensity. Sometimes this tendency is harnessed consciously: game developer Ian Bogost's 'procedural rhetoric' approach, for example, explicitly acknowledges and accounts for the fact that the developer has loaded their code with value judgements, in an attempt to change the user's behaviour. But far more often, such valuations are left implicit and unquestioned.

In years to come, this tendency will act as a brake on the spread and adoption of context-aware systems, if for no other reason than that the values and orderings encoded into a system at the time and place of its development will rarely prove to be the appropriate ones for other venues and audiences. Further, the nature of distributed, networked culture being what it is, it is a certainty that 'other' audiences will almost always outnumber those for whom a system was originally devised.

We've now spent a few thousand words problematising the more ambitious sorts of context-aware system, pointing out the deep issues to which they are subject, at least as most are presently constituted. At least for the time being, there would appear to be a certain hubris in daring to build educational strategies around them. Nevertheless, there is still room for genuine optimism here, albeit one significantly constrained by the limits of our tools and understandings. What problems in education, then, might location-based and



context-aware techniques actually be good for? How might they be designed with an appropriate degree of circumspection, so as to gain the greatest benefit from the available technology?

Games and other novel approaches to learning

One of the most promising aspects of location-based learning is that it might offer new hope to those for whom learning from books or in conventional classroom settings is difficult. Pervasive games and simulations, in particular, appear to offer educational content in a way that is commensurate with a wider range of learning styles.

Consider the Big Games developed by New York City-based design partnership area/code. These are physical challenges that, while mediated digitally, unfold against the terrain of the city; for example, players of the capture-the-flag game Crossroads (2006) are given GPS-equipped handsets featuring a five-block-square grid of Manhattan's West Village, on which their position is mapped and on which they compete not merely with another human player but with an entirely virtual opponent, the dread Baron Samedi of voodoo lore.

Substantiating the argument Kipfer makes above, at least in part, I will vouch that the experience of playing Crossroads is uncanny: one is simultaneously moving through real space and also through the representational space of the game, able for the first time to draw a visceral connection between what is done with the body and what happens on screen.¹ What is more, I underwent a full adrenal response to the approach of Samedi, who was no less terrifying or threatening for being nothing more than a few pixels bobbing across a 12cm screen.

It is easy to imagine historical education, particularly, being enhanced by such techniques: the ability to situate what is being taught in the actual physical context in which it occurred – for a student to reckon with his or her own body just how small in scale a medieval guild city was, for example, or to trace the tactical thrusts and parries of history's great (and not-so-great) battlefields on the very terrain involved – would clearly be invaluable.

Along these lines, Karen Schrier, a designer working for the US-based children's media network Nickelodeon, developed a game called 'Reliving the Revolution', or RtR, in 2006. RtR uses handheld devices to overlay information on the Revolutionary War battlefield of Lexington, Massachusetts, imposing roles and tasks on players to motivate them through the physical acquisition of what they may barely register as historical knowledge. (Note also how the shift from 'student' to 'player' works to do much the same here.)



¹ Another area/code-developed game inverts the equation: the 2007 Shark Runners, conceived for the Discovery Channel, is meant to be experienced in a relatively static environment, but builds its game play on real-time positional telemetry from sharks in the wild.

RtR is a just a first step, but there's little doubt that there is enormous promise in the fusion of visceral, Crossroads-style gameplay with an explicitly educational mission, beyond platitudes about 'learning made fun'. What's more, such 'augmented reality' techniques are already finding their place in the broader educational domain, most notably in vocational training for maintenance personnel. Employees of the Boeing aircraft company, for example, can push their head through an airliner's service hatch, survey the tangle of cabling confronting them and be offered a visual overlay explaining which wire belongs to which subsystem, which requires rerouting, and which must under no circumstances be cut.²

Such informational overlays extended to the terrain of a city, especially when opened up to user-generated content, will utterly transform the ways in which we teach, learn about, and, it must be said, understand the dynamics of structuration that govern the places we live in. Simply being able to construct 'heat maps' of affinity and use – see, for example, the house-by-house maps of real-estate valuation published by zillow.com – will surface and make explicit so much that has lain latent throughout the history of human habitation, allowing urban planners to make canner decisions about resource allocation. If nothing else, this newfound ability to record our subjectivities by the millionfold and moor them in place is a potential boon for future sociologists, who will enjoy an inconceivably better grasp of economic geography and patterns of urban use than we do now – learning at the broadest scale and of the deepest order.

New approaches to assessment

Somewhat closer to home, one area where a context awareness constructed with appropriate circumspection may actually prove useful is the domain of computer-adaptive assessment, or CAA. Indeed, much of the research in ubiquitous educational systems so far has been devoted to so-called 'portable assessment' tools.

The premise of CAA is simple, and is on its face hard to find fault with: tests are composed dynamically, in real time, in response to the test-taker's answers to previous questions. The difficulty of successive questions can be tuned to account for patterns beyond mere correctness – for example, the time interval required to complete a given question. The arguments that are advanced for CAA include that it ostensibly increases efficiency, reduces student anxiety, and provides a more accurate picture of the test-taker's true competency than traditional methods of assessment.



² Indeed, the real-time provision to the maintenance worker of entire databases full of vehicle schematics and the relevant repair instructions has blurred the distinction between on-the-job training and task execution. Under such circumstances, the risk of permanently deskilling the worker/learner is a concern that needs to be taken very seriously. As Marshall McLuhan would have reminded us, "Every extension is also an amputation."

It seems entirely plausible to suggest that the spread of cheap, robust environmental sensor grids will render other indices of performance trivially accessible to the testing authority. Particularly, easy-to-garner biometrics such as pulse and respiratory rates, temperature, pupil dilation and galvanic skin response will probably be used to evaluate not merely a student's level of competence, but their comfort with that competence.

On the upside, this presents the felicitous possibility that assessment can more closely come to resemble what it is supposed to be in theory, a tool harnessed to produce better educational outcomes rather than an end in itself. Anything that would help students gain an accurate overview of their own retention and understanding of knowledge, at the pace and in the mode most congenial to them, is to be welcomed. There are equally clear downside risks, however, not least that a progressively wider band of everyday experience will become subject to assessment, regulation and grading. Whatever benefits may be derived from context-aware CAA – a better acronym than 'CACAA' will clearly have to be found – it also, of course, speaks not at all to the curriculum involved, or to the content of what is being tested. This is of no small concern in an era in which, at least here in the United States, a plurality of constituents has apparently opted out of reason, science, critical thinking, and the methods associated with them.

Always new horizons

Finally, and much more stirring, the sort of systems we've been discussing offer a clear promise as a flexible, distributed support environment in which contemporary e-learning strategies can be exported from the desktop. The idea that anywhere can be a classroom is particularly interesting in the context of developing countries, but also of concern anywhere formal educational environments are absent, degraded, or dangerous.

It should be admitted that education will, of course, be competing against a great many other prerogatives for use of the local information infrastructure, and that for a variety of social, economic, emotional and psychological reasons, these other claims will often be seen as more pressing. Bearing in mind, as well, that the selfsame ubiquitous technology underpinning the promise of location-based learning also means that any room can be an office, a trading floor, a cinema or a sweatshop, it is still inspiring to think of what can be achieved at the intersection of place, social context and appropriately-designed systems.

What is at stake is nothing less than a reappraisal of what we mean when we say 'education'. Rather than something abstract and detached from the context which lends resonance and interest, certain kinds of knowledge can be re-imagined as a property of place itself, as something more akin to *genius loci* than to anything we'd recognise as a 'lesson'. Certainly, the teaching of disciplines as diverse as economics, history, physics, and anatomy will be

transformed by objects, transactions and places endowed with the ability to speak themselves – an ability inherent in almost all schemes for the deployment of ubiquitous informatics now being contemplated.

The success or failure of such initiatives will hinge to a great degree on decisions made at the level of their architecture, and to the humility and realism with which they are devised. The clear lesson of the last decade of research into the question is that the promise inherent in location-based and context-aware techniques is very real, but that delivering the promise – especially in the educational domain – will depend vitally on the degree of insight and sensitivity the designer is able to call upon.

As a long-time student of ubiquitous information-processing systems, their development and their use, it is my belief that it's a serious mistake to think of such systems as primarily technical in nature: their prospective deployment in the everyday, among literally billions of non-specialist users, brings them into the proper ambit of the social sciences. At the same time that Mark Weiser spoke of an information technology that was 'invisible, but in the woodwork everywhere', he also reminded us that 'the most profound technologies are those which disappear'. The responsibility incumbent on those who would design, develop, and deploy any such technology – and especially on those who would bring it to bear on questions of teaching and of learning – is self-evident.

Emerging trends in serious games and virtual worlds

Sara de Freitas [Serious Games Institute](#)

The role of 'serious games' in modern culture is a recent phenomenon, and broadly arises out of the wider use of electronic gaming for leisure purposes and the increasing use of the internet to support large online communities. Serious games, as distinct from leisure games, provide users and players with opportunities to explore non-leisure applications using games and immersive world applications for education and training, as well as supporting business and medical uses (Michael and Chen, 2006). The term has been coined to create a separation between leisure and non-leisure games-based activities in order to take games as training or learning tools more seriously. The use of serious games, in this way, may engage under-served learners, liven up school and tertiary curricula or provide support for lifelong learners in new and innovative ways.

The emergence of virtual world applications such as Second Life and ActiveWorlds provides potential for supporting learning communities in new ways. Virtual world applications, like immersive serious games applications, offer the capacity for using three-dimensional spaces as new learning spaces. This can support seminar activities, streaming lectures, create cyber-campus and help to support distributed and remotely located learner groups. This may add value to existing educational provision, as well as extending new provision of learning.

Serious games and virtual world applications offer great potential for learners to step inside the screen of their imagining with such possibilities as role-playing characters from history to re-enact events such as in the game mod (modification) of *Neverwinter Nights*, *Revolution*, which was modified by researchers at MIT in order to study the effectiveness of game-based learning with students of history. The idea was to help students to role-play social characters during the American Revolution to allow them to empathise with the people from that time. These formats can also be used to role-play researchers perhaps interviewing famous scientists long deceased, or as scientists undertaking experiments only possible in outer space (see Figure 1). In this latter example, students can use the tool developed by researchers at the University of Wisconsin-Madison to enact physics experiments that cannot be employed in the real lab. Learners could also become virtual tourists visiting museums a thousand miles away (de Freitas, 2006; Sandford *et al.*, 2006). In addition, these applications are supporting a whole host of social interactions providing scope for learners to meet with mentors and subject experts from around the world, undertake virtual work experience or form a distributed learning community to solve challenges and problems, play educational games and share and produce content.




Figure 1: Supercharged used for physics students.

Reproduced by kind permission of Kurt Squire



Figure 2: Second Life.

Image reproduced by kind permission of Linden Labs UK.

Serious games and virtual worlds allow us the potential to:

- provide support for our learning communities
- broaden our networks of learners
- provide tools to support creative learning activity and experience design.

Part of the problem serious games set out to address is the gulf between learners' experiences with technology inside and outside formal education. For some at least, this provides a real opportunity to extend learning beyond the conventional boundaries to the widest number, providing scope for reorganising learning and designing learning activities and interactions to fit infinite possibilities. The challenge that faces us today is how we can best make use of these applications to support learning.

The area of 'serious games' and virtual world applications therefore encompasses a wide range of applications including the following:

- Serious games applications such as Flash- and Java- based animations, immersive 3D single- or multi-player games developed on proprietary platforms (see Figure 3). One example of the latter type of serious game application is the demonstrator being developed by TruSim (a division of Blitz Games). The demonstrator mocks up an explosion in a busy urban area, the learner role-plays the medic arriving at the scene

of the explosion and is tasked with sorting through the casualties (Triage Sieve) in order of urgency. Increasingly, the crossover between leisure and non-leisure games is leading to more mainstream serious games such as Dr Kawashima's Brain Age and Big Brain Academy.

- Virtual world applications such as Second Life (see Figure 2) and ActiveWorlds, which are becoming popular with users – and more recently learners. Increasingly, universities are modelling their campuses into Second Life and onto other virtual world platforms. This virtual presence is allowing universities to reach new audiences and early adopters are utilising the capability to offer virtual seminars, streamed lectures and presentations. Virtual conferencing is a popular application with businesses, as well as recruitment, communications and marketing.



Figure 3: Triage Trainer.

Reproduced by kind permission of TruSim (A Division of Blitz Games)



Figure 4: Screen shot from emergency training session using Forterra. *Developed by Forterra Systems.*

One interesting example of how Second Life is being used is the use of Teen Second Life (for use by teenagers only) by the Open University. Here, students from the National Association for Gifted and Talented Youth (NAGTY) are taking lessons in virtual classrooms. The pilot developed under Schemes, is a project aimed at developing new education systems in both real and virtual worlds. [http://schome.open.ac.uk/wikiworks/index.php/The_schome-NAGTY_Teen_Second_Life_Pilot]

- Serious games applications that can be created in construction toolkits whereby users can build and develop their own scenarios and worlds, such as Olive on the Forterra platform (based upon There.com). An interesting example of this trend is the emergency training sessions led by Stanford University and being piloted using Forterra (see Figure 4). Previously Stanford was closed down to facilitate emergency training sessions involving different organisations including medical and police forces. For the last two years training has been taking place in the virtual world and this has allowed distributed teams to role-play serious incidents using real people to play act victims and for the trainees to undertake lifelike training. This form of application has real potential for training with different organisations and allows us to mock up an infinite range of scenarios.

The capability of the 3D web extends and enriches the potential of the 2D web (including existing tools such as FaceBook, MySpace, Bebo and Habbo Hotel). The possibilities of the 3D web include the capability to integrate these and other 2D tools and toolkits, to support learner-generated content, to enable sharing of content and to allow us to visualise more clearly different scenarios of practice. However, increasingly the move towards the 3D web is also prompting the potential of convergent technologies and applications, such as mobile gaming, alternate reality gaming and augmented reality gaming, which are allowing for bridging between real and virtual spaces. This trend is having an impact upon design in the real as well as virtual worlds, with for example increasingly flexible designs of spaces being used in the real world to reflect greater possibilities and tie-ins with the virtual world experience.

The issue of definitions is one that has been and will continue to be hotly debated. Table 1 gives an idea of the different definitions and terms that are emerging daily in the field and can distract or put people off engaging with the area. An added complication is explaining the range of different kinds of applications available, and while it can be easy to show clips from demonstrators to illustrate this, those not familiar with the levels of fidelity and interaction can find it hard to imagine the possible uses.

Table 1: Definitions and terms of games

Term used	Related or synonymous terms	Descriptions & references
Educational games	Computer games, video games, serious games, game-based learning, instructional games	Games in general can be defined in surprisingly numerous ways, often changing the way games are used and perceived (Wittgenstein, 1958). Games are often defined as 'a series of choices' or as 'rule-based play'. For the purposes of this report, educational games for learning, like serious games, are defined as: applications using the characteristics of video and computer games to create engaging and immersive learning experiences for delivering specified learning goals, outcomes and <i>experiences</i> .
Online games	Massively Multiplayer Online Role play Games (MMORPGs), Massively Multiplayer Online Games (MMOGs), persistent games, MMORTS (Massively multiplayer online real-time strategy), MMORPGs (Massively multiplayer online first-person shooter)	Online games are becoming more widely used since their emergence as multi-user dungeons / dimensions (MUDs) in the 1980s. Online games include simple text-based games as well as games that involve complex graphics and virtual worlds that are used by large numbers of players simultaneously. Broadband access to internet resources has made MMORPGs, MMORTS and MMOFPS very popular. In addition, the wider usage of Flash and Java has allowed gaming websites to use streaming video, audio, and introduce greater user interactivity.

Term used	Related or synonymous terms	Descriptions & references
Serious games	Educational games, video games, game-based learning, instructional games, sim games, gamesims	Michael and Chen (2006) give the following definition: 'A serious game is a game in which education (in its various forms) is the primary goal, rather than entertainment'. The following definition has been developed for the SG-ETS project: 'Serious games for learning are applications using the characteristics of video games to create educational and engaging learning experiences and deliver specified learning goals.' It is worth noting that Huizinga defined games as a free activity standing quite consciously outside 'ordinary life', as being 'not serious' (1980), and following this definition games cannot be serious. Callois similarly defined games as voluntary and therefore also conflicts with the notion of serious games (1961: 10-11). This gives a good indication of the kinds of contradictions found in comparisons of the available literature.
Simulations	Electronic simulations, virtual reality systems, training simulations or simulators	A computer simulation is a way of modelling a real-world situation on a computer. By altering variables, predictions about the behaviour of the system may be made. Simulations have traditionally been considered as types of games. But equally the earliest simulations were war games. The relationship between games and simulations has been close, and even when virtual reality systems were being pioneered, the power of immersive environments for learning was recognised. Also, simulations may be defined as non-linear exploratory environments (Aldrich, 2004, 2006).

Term used	Related or synonymous terms	Descriptions & references
Serious virtual worlds	Immersive worlds, 3D environments, virtual worlds, virtual environments, metaverses	<p>The predecessors to the virtual world were the multi-user dungeons (MUDs) of the 80s, which had all the characteristics of the modern virtual worlds but were text-based. These media forms provided the foundations for the development of online communities supported in their daily activities by 3D and animated spaces providing a backdrop for the day-to-day activities that take place there. A serious virtual world is an environment where players and users can explore a 2D or 3D world, freely taking on the identity of an avatar (which represents the player in the virtual world), play games and participate with online communities. The earliest virtual world with avatars dates back to LucasFilms' Habitat in 1985 (Morningstar and Farmer, 1993). Over the last five years the use of virtual worlds for educational purposes has grown, including replicating universities, museums and art galleries, and science labs to creating fictional worlds for tutoring and mentoring (Prasolova-Førland <i>et al.</i>, 2006). Most of the main open-ended virtual worlds such as Habbo Hotel, ActiveWorlds and Second Life use avatars, allow the creation of objects and construction of buildings. Some virtual world applications such as Guild Wars utilise a narrative and have quests, users can join guilds and fight monsters or collect objects.</p>

The debate about definitions perhaps reaches its apotheosis in the arguments of Seymour Papert who debates the notion of 'edutainment' as neither as engaging as leisure games nor educational (1998). The hybrid form of edutainment like serious games does for some at least point to an anomaly: how can games be educational, how can they even be serious? In addition, many tutors and practitioners are keen to find out how effective game-based learning is and how they set about getting involved with game-based and virtual world learning activities.

Trends

The social and the immersive

The main change agents in the current environment can perhaps be posited as the wider uptake of the internet and globalisation. This has a social implication that is undeniably changing our social structures – including schools, colleges and universities – as well as substantial changes upon continuing professional development (CPD) and work-based learning.

The trend towards self-organising communities, flatter social hierarchies and the potential for distributed activities are having a reinforcing impact upon globalisation, making the world smaller and the interconnections and scope for social interactions greater. The drive towards greater potential for social collaboration, through social networks and often web-based self-organising communities, has led to sets of tools that are encouraging a significant trend towards mass user generated content (at present a quarter of all data is original) and IDC estimate that 70 per cent of content by 2010 will be user generated, on top of growing participation levels that were at 1 billion in 2006. The second result of the shifts of globalisation and the internet can be felt in the drive towards the 3D web, a new vision for information access and use centring upon visual rather than textual data, where the opportunity for user generation is key and where the integration of different media and tools is also significant.

Many talk of the 3D web and expect its use to grow significantly. In particular the trend amongst large multi-national companies such as IBM, PA Consulting and Reuters is towards exploring possibilities for collaboration and communication. Using 3D interfaces to a range of resources, materials and communications is a clear trend, with the emergence of 'mash-ups' where different programs are brought together, such as Google maps and databases with information about a local area. The trend is supported by the easier integration of programs using open standards and interoperability, and there is now a move towards interoperability standards for virtual worlds (see: http://www.news.com/Tech-titans-seek-virtual-world-interoperability/2100-1043_3-6213148.html). The role of distributed social networking is also becoming a clear trend with businesses notably taking up the new media technologies to support business applications and communications, as well as supporting applications from training to recruitment and providing a focus for interactions with customers.

Converging technologies, diverging applications

The greater scope of the convergence of 3D technologies, applications and tools with other media supports and extends this trend towards supporting evolving and self-organising communities. The rate of convergence of games technologies with other technologies is significant. For example, recent work

being undertaken is looking to build bridges between these different technologies. Alongside converging technologies is the trend of diverging applications from business solutions, to training, recruitment to work experience, sales and marketing to communications within and outside of the company and the range of emerging applications is diverse. This trend of converging technologies and diverging applications is unprecedented, and it is difficult to think of a media form that has led to such pervasive uses and applications as we have noted with the internet and game-based and virtual world applications.

Another strength of these applications is the capacity for integrating with different media and interactive resources, as well as integrating with available social software and collaborative tools such as live chat facilities, bulletin boards and shared resources, which means that these applications have real potential for supporting distributed communities in different geographical locations, or special interest groups, or mixed groups of learners (age, gender, nationality etc.). While for standard tutorial group structures this may appear less appealing, it does open up the option for learner groups studying out of normal hours of learning, or beyond the timeframe of the course, and open up real potential for learning outside the standard institutional framework.

The fast changing converging technologies offer a real challenge for IT support, particularly where a culture of a closed system has been adopted for security reasons. However, a balance between security and open access needs to be found. Institutions need to provide ways of supporting new emerging technologies and applications and this can be difficult in terms of staff training as well as needing a more flexible approach to IT support. The moves towards service-orientated architecture and interoperability may help to lessen some of these challenges and remote IT support is becoming more possible, but the role of the institution clearly needs to find new strategies for adapting to the proliferating 'worlds' particularly if it is to continue to remain relevant to those learners regularly engaged in exploratory spaces, using games applications and *au fait* with a range of different social software tools.

Sharing learner-generated content is becoming a way of life now, and the role of education is to take up these tools and help learners to become more adroit at using them, rather than ignoring or banning them. As practitioners this may indeed become a creative process rather than one to be feared. The new technologies offer the learners a chance to take a more empowered role in the learning interactions and activities, but also allow for deeper reflection upon learning, sharing learning resources and outputs and engaging under-served learners in new ways. In addition, the critical role of authoring content offers a diverse range of new approaches both to design and to learning. However, curricula and institutional structures may need to loosen up to adapt to the rapid change, and it will be important for tutors and tutor practitioners, as well as policy developers and senior managers, to make sure that the critical skills needed to

remain analytical are maintained at the centre of games and activity development and usage. The danger of using so many different communication channels simultaneously is that attention and focus may be difficult to hold, so work to enforce academic rigour, analysis and synthesis, as well as meta-reflection and higher order cognition, needs to be considered in learning design. The drive towards a more seamless learning experience more focused upon social interactions would benefit the learner through enriching learning, connecting with real experiences and reinforcing learning through social interactions.

Case study: Teen Second Life

The Open University undertook a study to evaluate the educational potential and pitfalls of Teen Second Life. The study aimed to find out more about how 14–19-years-olds use Second Life, and used 149 students from the National Association of Gifted and Talented Youth (NAGTY) as a sample. The study found that the level of engagement was comparable to other media. Of the 68 per cent of students visiting the island, 41 per cent spent more than an hour. Access was an issue particularly for some students. But those who could access the Second Life site Schome Park (<http://www.schome.ac.uk>) and participated with the wiki and forum developed a wide range of Second Life skills (such as moving around the environment, scripting and producing movies in-world and constructing objects), skills which were found to have transferable value. In particular the study highlighted communication, teamwork, leadership and creativity as the particular skills supported through the activities. Interestingly, those who used the wiki and forum showed higher levels of performance than those who just participated in Schome Park. Notably, students that found social aspects of life problematic at school found that Schome Park provided a more secure and safer environment within which to explore social relationships, leading to enhanced confidence and the development of social skills.

Quote from student:

'I think that what Schome is doing through breaking down the barriers between teachers and students making it hard to see where one stops and the other begins, is fantastic, because when everyone is on the learning curve together, it brings about less of a feeling of segregation and a greater feeling of equality, and this leads to trusting people more...'

(Twining, 2007).

Issues

Research overview

One of the issues with research in the field has been the lag between findings and innovative practice, and while this problem is not a new one it seems to be exacerbated in the current climate of fast convergence and rapid innovating practice. While much of the evidence for serious games, like e-learning in general shows 'no significant difference' when compared with face-to-face, blended modes of learning, when learners use multi-modes of learning, often accelerated learning and longer retention of information results.

The efficacy of games in studies, like e-learning in general, has been inconclusive, with 'no significant difference' being reported in some studies where face-to-face and game-based approaches are set head to head. However, as with other media-based learning, most evidence has pointed to blended approaches to learning being more effective than one or another, and this needs to be reflected in experiment design. Some data has come from surveys with users, and these studies have in general shown that some learners do not like using game-based approaches (de Freitas *et al.*, 2006), but evidence from other studies indicates that this form of learning can change attitude (Hays, 2005) and that it can be engaging and motivating for learners (Garris *et al.*, 2002; Mitchell and Savill-Smith, 2005), as well as being helpful for engaging under-served learners (de Freitas *et al.*, 2006).

However, for certain user groups, in particular under-served, more visual and younger learners (see for example de Freitas *et al.*, 2006), games have been found to have positive benefits. Moreover, when users are part of the design of the games (in particular using participatory design methods) the games may be substantially better taken up and are often more effective following as they do the learners' needs and requirements (see for example Dickey, 2005; de Freitas and Oliver, 2006). Despite some negative psychological studies finding a link with leisure game-play and addiction, the use of educational games and simulations on the whole does not show addictive behaviours, possibly because they are not as engaging as leisure games (de Freitas and Griffiths, in press).

While leisure games studies have shown that age, gender and cultural differences pervade, a recent study is revealing that there was no significant difference ($p < 0.05$) between age, gender and culture regarding the use of specifically educational games. The finding confirms that serious and leisure games are still distinct categories, also indicating that educational games may be used with mixed age, gender and cultural groups to equal effect, as supported by other studies (de Freitas *et al.*, 2006). The study which surveyed medics (nurses) also revealed that the 40+ group played significantly fewer games than the <30 and 30-39 age groups, and that the patterns of game-play between <30 and 30-39 were not as significantly different as would be expected.

An interesting confirmation of anecdotal evidence is emerging, which indicates the power of the format for engaging learners particularly in the under-40 age group, but studies have not yet confirmed the full power of educational games for older learners (preliminary findings as presented by de Freitas, 2007; Jarvis *et al.*, 2007; data to be published).

Potential uses

Clearly, ‘immersive world’ applications have the potential to support communications between learners, to support problem-based learning opportunity and to support exploratory learning experiences (Saunders, 2007). However, much needs to be understood about how to best convert these spaces for learning purposes such as seminars, simulations, modelling, learning activities, networked learning experiences, cybercampuses and streamed lectures (Prasolova-Førland *et al.*, 2006). While the spaces are excellent for bringing together the use of a range of different media (streamed video and audio, email, live chat, social network software, mind mapping software and others), questions remain as to how best to integrate these media to support the most enriched learning experiences.

Ohio State University is one example where Second Life is being used innovatively to support enriched learning experiences. The presence includes three units: OHIO Outreach, Ohio University and Ohio STEAM on the teen grid. The model centres upon serving traditional university students, distance and adult learners and high school students in a way that supports both live and asynchronous learning experiences. Using both futurist spaces and models of real-world buildings, the campus aims to engage learners, develop a unique identity for the university and to integrate a range of teaching tools such as voting and survey tools, learning objects and spaces for seminars and lectures. (See the video clip from Ohio State University campus at <http://youtube.com/watch?v=aFuNFRie8wA>)

The research has revealed that many of the early adopter groups have been in vocational training areas, which is not unexpected as games and immersive worlds applications have real potential with respect to experiential and exploratory learning models. In particular, the issue of learning transfer, which lies at the heart of the debate about the efficacy of game- and simulation-based approaches, relies upon a degree of fidelity to the real work experience, and here games technologies are making leaps and bounds. Cost has also been a factor in the pattern of uptake in the new sector, and areas with large numbers of students, or where training has life and death consequences (such as medical and military applications) have reflected this imperative. However, with the emergence of immersive world applications, such as Second Life and ActiveWorlds, this trend is changing and new areas of training involving cross-agency training, emergency training and professional training across different areas are becoming possible. The fact that new areas of training are now



possible through these new means implies a greater need to consider the pedagogic underpinning of training in virtual worlds, and frameworks and models (such as Four Dimensional Framework and Becta's Quality Principles for Digital Learning Resources) are slowly emerging to help us to design, develop, select, use and evaluate serious games and virtual worlds applications.

One example of medical use of serious games demonstrators is the Pulse!! project – The Virtual Clinical Learning Lab. This is a virtual training environment designed to support a range of the training needs required by nurses and medical professionals. The US Department of the Navy's Office of Naval Research is funding the immersive virtual learning space, being developed by BreakAway Ltd, for the Texas A and M University-Corpus Christi. Virtual patients, using artificial intelligence (AI) will respond in lifelike ways to environmental changes and medical techniques and skills used by the trainees. The system may be used by new trainees or for established professionals to update training. (Johnston, 2007; see also www.sp.tarnucc.edu/pulse/index.shtml).

Alongside CPD and training, and cross-agency distributed learning opportunities, scope for learning in colleges and universities using seminars in SL, sharing resources, cross-disciplinary teaching, team teaching and wider use of mentoring are emerging. Aspects that may well help to enrich the learner's experience and allow them scope to become more active in their learning, through developing their own scenarios of practice, their own content and their own perspectives upon reflection and meta-reflection. The scope for virtual work experience, for example, would allow learners to create better links between their studies and their chosen path in life.

In the field of science education a number of exciting tools and games are emerging to bring to life experiments that in some cases could not be undertaken in the lab. The scope for games and immersive world applications to reach under-served learners, and learners with no previous interest in science education, could unlock a new enthusiasm for subject areas where traditionally few learners have participated (de Freitas *et al.*, 2006). The tools also allow for cross-disciplinary teaching, which could also allow tutors to engage learners in weaker areas.

An interesting example is SciLands, the science and technology region of Second Life. The region is well worth a visit and includes many highlights including the International Spaceflight Museum and the recently launched Nanotechnology Island. Resources available to this community include, shared resources for those in formal and informal science education and an extended network of specialists and organisations. Imperial College, London and the National Physical Laboratory have developed Second Health, a 3D vision that provides a vision of health of the future (see the video link to Second Health: Polyclinic Tour, at <http://secondhealth.wordpress.com/movies>).

The potential uses of games, simulations and virtual worlds are many and varied and it will be interesting to see how these are exploited within academic research and teaching circles over the next five years. Will the emphasis upon collaborative learning, for example, lead to new teaching and learning models, will learners more generally become the 'authors' of their own learning materials? Will tutors become the choreographers of learning experiences through designing activities and scenarios, and through mentoring?

Impact upon learning

The use and adoption of virtual spaces has an implication upon where, what and how we learn. This change is already being felt, as we design our physical spaces to reflect the need for multi-usage of space, and as learning itself becomes less about conveying information and more about designing experiences and activities. The emphasis upon more seamless learning experiences that may cross the boundaries of real, virtual or imagined spaces, has challenges for us of course, as managers, as tutors, as learners and as designers or choreographers of our own learning trails or pathways. In the end, it would seem as though learning becomes less about the ability to reproduce standardised components of learning, and more about allowing individuals to inform and design their own learning interactions and transactions.

This possibly more empowered model of learning not only challenges us but also, and centrally, our institutions and organisations, favouring more decentralised structures with more dynamic and adaptable structures and self-organising communities emerging to support change, not like the current personalising learning agenda.


The impact of games and immersive world applications upon our learning therefore may be substantial, and for this reason needs to be explored by managers, researchers, tutor practitioners and learners. The need to produce guidelines, supports and communities to help us both understand and research new learning experiences is critical to the ongoing development of our online and offline communities. However, we need to adopt the same critical approaches that we apply to the real world in order to allow us to reflect upon and become part of the design mechanisms emerging to support the wide uses and applications of games and immersive worlds.

Collaborative learning is definitely not a new prospect for educators and learners, as most learning takes place collaboratively in groups as part of social interactions. However, the major paradigmatic changes occurring with globalisation and the use of the internet are making our worlds smaller and allowing us to reconsider how we learn, where we learn and what we learn. Perhaps what we are really talking about is the need to adapt to a changing vision for education based rather upon more immersive learning experiences and social interactions.

References

- de Freitas, S. (2006), *Learning in Immersive Worlds*. Bristol. Joint Information Systems Committee. Last retrieved online 1st November 2007 at: www.jisc.ac.uk/eli_outcomes.html
- de Freitas, S. and Griffiths, M. D. (In press), 'Massively multiplayer online role-play games to support learning and training communities' in R. Ferdig (Ed.), *Handbook of Game-based learning*. Idea Group Publishing.
- de Freitas, S. and Oliver, M. (2006), 'How can exploratory learning with games and simulations within the curriculum be most effectively evaluated?' *Computers and Education*, (2006) 249–264.
- de Freitas, S. (2007), 'Building effective exploratory learning experiences with serious games applications', Presented in May 2007. Vienna Subotron Lectures, Vienna, Austria.
- Dickey, M. D. (2005), 'Engaging by design: how engagement strategies in popular computer and video games can inform instructional design'. *Education Training Research and Development*, 53(2): 67–83.
- Garris, R., Ahlers, R. and Driskell, J. (2002), 'Games, motivation and learning: a research and practice model'. *Simulation and Gaming*, 33: 441–467.
- Hays, R. T. (2005), *The effectiveness of instructional games: A literature review and discussion*. Technical Report. Naval Air Warfare Center Training Systems Division. Orlando, Florida.
- Jarvis, S., Hallam, C. & Smith, S. (2007) Preliminary findings from research on the causes of failure to adhere to infection control policy in two NHS Trusts. Health Protection 2007, 17th September.
- Johnston, C. (2007), 'Pulse!!! The Virtual Clinical Learning Lab'. Human Factors Integration Defence Technology Centre Newsletter, 7: 4–5. Last retrieved on 21st August 2007 at: www.hfidtc.com
- Michael, D. and Chen, S. (2006), *Serious games: games that educate, train and inform*. Boston, MA. Thomson Course Technology.
- Morningstar, C. and Farmer, F. R. (1993), 'The Lessons of Lucasfilms' Habitat' in M. Benedikt (Ed), *Cyberspace: First Steps*, pp. 273–301. Massachusetts. MIT Press.
- Papert, S. (1998), 'Does easy do it?' *Games Developer*. June. Last retrieved online on 16th October 2007 at <http://www.papert.org/articles/Doeseasydoit.html>
- Prasolova-Førland, E, Sourin, A. and Sourina, O. (2006), 'Cybercampuses: design issues and future directions'. *Visual Computing*, 22(12): 1015–1028.
- Saunders, R. L. (2007), 'The genesis of a virtual world revisited'. *International Journal of Web-Based Communities*, 3(3): 271–282.

Selected resources

-  de Freitas, S., Savill-Smith, C. and Attewell, J. (2006), *Educational games and simulations: Case Studies from Adult Learning Practice*. London. Learning and Skills Research Centre.
- Egeneldt-Nielsen, S. (2005), *Beyond Edutainment: Exploring the educational potential of computer games*. Doctoral Thesis, IT-University of Copenhagen. Denmark.
- Entertainment and Leisure Software Publishers Association (2006), *Unlimited Learning: Computer and video games in the learning landscape*. London. Entertainment and Leisure Software Publishers Association.
- Gee, J. P. (2003), *What videogames have to teach us about learning and literacy*. New York: Palgrave Macmillan.
- Mitchell, A. and Savill-Smith, C. (2005), *The use of computer and video games for learning. A review of the literature*. London. Learning and Skills Development Agency.
- Salen, K. and Zimmerman, E. (2004), *Rules of play*. Cambridge, Massachusetts. MIT Press.
- Sandford, R., and Ulicsak, M., Facer, K. and Rudd, T. (2006), *Teaching with Games: Using commercial off-the-shelf computer games in formal education*. Bristol: Futurelab.
- Twining, P., *et al.* (2007). The schome-NAGTY Teen Second Life Pilot. Final Report. Open University. Retrieved online, 14th November 2007 at: <http://kn.open.ac.uk/public/getfile.cfm?documentfileid=11344>
- Williamson, B. and Sandford, R. (2005), *Games and learning: A handbook from Futurelab*. Bristol: Futurelab. Retrieved online, 14th November 2007 at: http://www.futurelab.org.uk/resources/publications_reports_articles/handbooks/Handbook133

'If it quacks like a duck...'

Developments in search technologies

Emma Tonkin UKOLN

Digital content is widely viewed as a viable replacement for traditional educational resources. Projects such as One Laptop per Child place interaction with digital content directly in the centre of a constructivist approach to education. There is a widespread sentiment that 'Google generation' children are sufficiently curious to educate themselves, given the ability to access information. How true is this? When we look at any new technology, we are usually too optimistic. Much as the internet itself in its early days was seen as a social enabler and a means to democratise society, the means to search and retrieve information is sometimes seen as a means to make knowledge available and therefore bypassing the intermediate steps of education.

In practice, however, there are complications, which can be identified by examining the capabilities and limitations of the technologies in question. There is a growing awareness that the high rank of a search engine result does not necessarily reflect the quality of the resource, and increasing concern that many are too quick to trust the ever-increasing information retrieved from the internet.



Traditional search engines are not all-powerful. For example, Google assembles evidence via links (the famous PageRank algorithm) and page content to index pages effectively. As such it is limited to pages that are *accessible* to its web spider, but many are not, giving rise to the *grey web* – publicly available pages that cannot be indexed. For the search engine words are simply abstract symbols with no underlying network of concepts, giving rise to a current interest in *conceptual*, *contextual* and *semantic* search, where pages are mapped against a structured set of terms or concepts. Multimedia content presents an additional set of practical difficulties, for the pragmatic reason that video streams are more difficult to de-construct into unitary features than a page full of text.

Search engines are not one-size-fits-all. Users vary greatly in needs and abilities through age, background, interests and task. Hardware form factors, abilities and interface peculiarities are significant factors in designing appropriate interface query modes. Sometimes a search is not what is needed at all, but discussion, browsing, or informal interaction with a peer group. A vast array of methods and technologies exist that, separately or in combination, allow search engines to respond to users' individual needs and circumstances. In this article we explore some factors involved in personalisation of search results, and several means of organising and searching through information: the Semantic Web; social tagging; web, multimedia and mobile search.

Digital objects without semantic annotation are just ones and zeros on a hard disk, just as a library without a catalogue is merely a stack of books. Some information is needed to enable the computer to locate objects usefully, but how much detail is required? Is this information extracted from the object or contributed by a human? Technical concerns abound; how should the data be encoded, and an object *annotated* with an accessible representation – *metadata*, data about data, relevant facts such as author, title and publishing date, or *relations* to other objects? Various mechanisms exist that provide a structured means of semantic annotation, including metadata standards such as Dublin Core (DC) – and the Semantic Web (SW).

The Semantic Web

Laying aside technical details, a computer needs some information about an object before it can recognise it as the resource that we are looking for, from the very obvious ('it's an image') to the detailed ('it's a picture of Alexander Graham Bell, who invented the telephone'). There is a reference problem associated to describing a digital object in terms of a number of entities and objects, 'telephone' and 'Alexander Graham Bell'; what are these terms? What do they mean? How are they related?

The set of concepts and technologies collectively referred to as the *semantic web* originated in 2001 with a landmark article written by Tim Berners-Lee, James Hendler and Ora Lassila (Berners-Lee *et al.*, 2001). Instead of providing unstructured records and then describing them separately, or extracting semantics from the object itself, the semantic web holds its own metadata. Data is published in a machine-readable form, so that a computer can interpret and apply the information available on the Web. The computer can then perform sophisticated tasks for the user.

For example, the computer, which lacks the commonplace knowledge required to reason about Alexander Graham Bell and the telephone, learns to relate concepts using well-formed facts, entities and relationships in order to reason about the world:

Alexander_Graham_Bell (is_a) man

A man (type_of) person

Alexander_Graham_Bell (inventor_of) telephone

telephone (type_of) communications_technology

The computer responds to queries about the invention of communications technologies by offering a resource depicting Alexander Graham Bell. Perhaps if it has enough records, it will also mention Elisha Gray, whose patent application arrived just after Bell's.

Berners-Lee's vision included more practical tasks. As such, the example scenario described two siblings, Pete and Lucy, arranging physical therapy sessions and chauffeuring duties for their mother's prescribed treatment. Instead of doing the research and scheduling the tasks themselves, Pete and Lucy each task their software agents (software that acts in the interests of a user, acting as their agent in a transaction) with researching the various *providers* who are able to offer the *prescribed treatment*, which are *in-plan* (within the expense budget), within a *20-mile radius* of their mother's home, and which are rated as *excellent* or *very good*. The agents then negotiate *appointment times*. Terms italicised in this paragraph are recognised by the agent as semantically meaningful.

Those who have chosen a doctor or healthcare provider for themselves or a dependent may not recognise themselves in Pete and Lucy. What are we really willing to automate based on a ranking score? One answer is that the process is not fully automated. The device retrieves recommendations based on a semantically marked-up subset of the data available on the internet about the various topics of interest. Moreover, if Pete or Lucy wish to check any review text that might accompany the various recommendations, they are welcome to

do so – but *sentiment analysis* of a text is a notoriously difficult natural-language programming problem, so they would be unwise to depend entirely on their agents' judgement in that area.

To a cynic, this scenario demonstrates nothing more than can be done using a set of databases. Travel agents have been able to recommend holidays according to distributed database records for many years now. This is true, but the point here is that arranging public access to those databases is not required. The data from these databases is published to the Web, where Lucy's agent can find it and use it as a basis for reasoning.

The Semantic Web in practice

The discussion surrounding the Semantic Web (SW) is one of extremes; a somewhat illusive, if keenly felt, dichotomy has been established between those who favour structure and those who feel, as Wittgenstein did in his later years, that concepts defy wholly objective description. Discourse is heated and often bitter, a warning sign that something far older and closer to the heart than technical concerns motivates the debate. The SW is frequently heralded as the dawning of a new era and has received substantial funding. Yet although a great deal of research is published every year, our talkative cynic might remark upon limited practical implementation outside the research community.

Like many ambitious projects in science, the Semantic Web has many detractors; as with superstring theory, there are few results that showcase the utility of the approach relative to alternatives. One high-profile critic, Clay Shirky, caused a flood of indignant responses with his 2003 article (Shirky, 2003), that described the SW as 'a machine for creating syllogisms'. This was inflammatory. Indignant commentators pointed out technical flaws, such as the point that the SW is largely based on first-order predicate logic (FOL), not syllogisms. Areas of life in which such logic is widely and successfully applied were held up to view. Yet Shirky's criticisms had many valid components. For example, his frustration with the level of complexity of the standards has been widely echoed by developers, and it is certainly true that there exist aspects of human behaviour and understanding that are not easily modelled in logic programming.

A quantitative analysis of the amount of Semantic Web pages on the Web from 2005 found that the majority of data available in the flagship standard of the effort, RDF (Resource Description Framework) is simply data encoded in RSS, the newsfeed format variously defined as Rich Site Summary and Really Simple Syndication, amongst others (60 per cent – versions 1.0 and 0.90 of the RSS standard are based on drafts of the RDF standard). A small amount is FOAF, or Friend Of A Friend, another technology based around RDF that was poised to take over the world in '03; social network applications such as LiveJournal, Facebook, Typepad and MySpace are the inheritors of the crown. Some make



use, incidentally or otherwise, of FOAF, but it seems that centralised storage of data has sufficed for social networking applications up to the present date. Only a minority of RDF records online represented novel SW developments.

Creation of a reasonably accurate model of actual experience – as opposed to an idealised model – is a difficult task. To quote McCool (2005), the ontological data model makes representation of any nontrivial factual information difficult because it can't represent context. Artefacts such as social networks (Paolillo *et al.*, 2005) are rather imprecise compared to the precise, static conception of meaning encoded into an ontology. Like any database, aging structured data needs maintenance. It is difficult to reengineer a system. Instead, we may resign ourselves to using an increasingly clunky and unintuitive system on the principle that it is a) familiar and b) still works.

As a rule-of-thumb, ontologies work well when they represent commonly held theories or models that are explicitly relevant to a given topic. For example, we apply a Linnaean taxonomy to categorising species, prioritising uniformity of classification over personal experience. One might find a platypus funny and a funnel-web spider horrifying, but subjective reactions do not form part of our formal classification of species and hence are understandably excluded from scientific discourse on the topic of Australian fauna. However, such nuances can be collected and added to our reasoning using other technologies, such as social tagging.

The role of AI and data mining

Who will generate SW annotations? There may be a tipping point. If organisation A makes its information available, perhaps organisation B will do the same, and eventually a critical mass of data will be met. But Berners-Lee expects the SW to be populated in part by means of automated approaches to information extraction from digital objects.

By 'semantic', Berners-Lee means nothing more than 'machine processable'. The choice of nomenclature is a primary cause of confusion on both sides of the debate. It is unfortunate that the effort was not named 'the machine-processable web' instead. This, along with some optimistic usage scenarios, is a primary cause of extraordinarily high expectations in those who take the term at face value, and hence to fierce criticism of what is seen as an overly-ambitious area of research. In summary, the SW is 'a webby way to link data'.¹ Those evaluating the technology are well advised to look at what it really represents today, rather than what it may one day become.

Initially, the SW was often presented as all or nothing, revolutionary technology. Who can blame commentators for seeing the Emperor's new clothes in such an ambitious technology? At present, practical deployment remains fairly 'heavy' and complex, involving an investment that, though rewarded with a more explicit and powerful representation, is arguably not always necessary. Much of the software remains somewhat experimental, and the IT workforce has limited familiarity with specific concepts or technologies.

Many see a pragmatic need for a compromise, a *slightly semantic* Web. Later, we will examine a few technologies, recently proposed or resurgent, that form part of today's functional compromise.

The lower-case semantic web

Microformats

Accepting that, for some purposes, the Semantic Web may be too much of a good thing, various forms of lightweight semantic tagging have been suggested. For example, the microformat was developed as a simple method of making Web pages that are a 'little bit semantic'. Appropriate attributes are placed in XHTML in order to render everyday information machine-readable. Current availability of marked-up content and services suggests an encouraging future (Allsopp, 2006), though it is too early to say whether the approach will be widely adopted in the long term.

Figure 1: Hcard allows us to easily mark up human-readable text; for example, UKOLN's contact details:

<pre><h2>Contacting UKOLN</h2> tel: ++44 (0) 1225 386580
 fax: ++44 (0) 1225 386838
 email: ukoln@ukoln.ac.uk web: www.ukoln.ac.uk </p></pre>	<pre><div class="vcard"> <h2>Contacting UKOLN </h2> tel: ++44 (0) 1225 386580
 fax: ++44 (0) 1225 386838
 email: ukoln@ukoln.ac.uk web: www.ukoln.ac.uk </p></pre>
--	--


A microformat reader recognises this as a visiting-card, which could be stored in an address book. This minor change for the author of the web page allows the browser to recognise the data and how it may be treated, stored and queried. Microformats exist for a few popular formats (addresses, event descriptions and news items, for example), but much of the text that is placed on the Web ‘resists simplistic representations’ (McCool, 2005).

Intrinsic and extrinsic meaning

Another popular technology today is social tagging, which is much less about establishing what a data object intrinsically is, and much more about finding out what people think about it and what it means to them. Social tagging systems (Guy and Tonkin, 2005) allow users to apply short plain-text descriptions to a data object. Tagging systems contain elements designed for personal information management (Kipp, 2006), topic-based keywords that would not be out of place in a controlled vocabulary, and informal descriptions. Tags may be applied by an individual for their own use or for a community or global audience (Tonkin *et al.*, 2008) and can be thought of as digital annotations in the sense described by Marshall (1997). Several widely publicised and popular sites make extensive use of tags, including Flickr [<http://www.flickr.com>], in which users tag their photographs, and del.icio.us, a social bookmarking service in which any user may tag and share any resource available on the Web. Tag systems may also contain other types of information, such as geotags (referent location information).

LibraryThing [<http://www.librarything.com>] is a great example. Despite its use of constrained subject vocabulary, LibraryThing, like most uses of social tagging, is all about aggregating points of view. On LibraryThing, Library of Congress Subject Headings (LCSH) are applied for book categorisation, supplemented by additional LCSH categories provided by libraries across the United States – and users’ own opinions. The site allows users to click through and see book recommendations and user preferences. The downside of this connectivity is the ease with which young readers may find themselves outside their comfort zone, an effect common to any highly connected network.

The difference between intrinsic and perceived meaning in recent library history can be illustrated by Ray Bradbury’s statement that his famous novel, *Fahrenheit 451*, is ‘not about censorship’ (Johnston, 2007). LibraryThing [<http://www.librarything.com/work/4248>] suggests that the consensus of opinion does not support Bradbury’s opinion of his work – the tag cloud for the book (Figure 2) shows the terms most commonly applied to the resource. Tag clouds are a common visualisation method for tags, and give a quick visual idea of the sort of terms applied to an object.

Fig 2: The LibraryThing tag cloud for *Fahrenheit 451*

(42) 20th century(58) american(100) american literature(72) book burning(75) **Books(111)** books about books(16) Bradbury(46)
 censorship(234) classic(342) **Classics(171)** dystopia(512) fantasy(44) favorite(24) favorites(17)
Fiction(1,091) fire(15) future(61) futuristic(29) high school(19) literature(141) novel(155) own(61) owned(25) paperback(45)
 political(24) Political fiction(18) politics(31) Ray Bradbury(51) read(226) reading(15) satire(25) School(25) **Sci-fi(348)** science
 fiction(971) sf(143) sf(15) Social commentary(43) society(17) Speculative Fiction(30) ste(16) totalitarianism(29) unread(48)

Tagging is widely criticised as ‘noisy’, and the wide range of annotations – in particular intended audience and level of formality – attracts criticism from commentators. Tags cannot be trusted; they are explicitly points of view. On the other hand, tagging systems often form part of a social network (Tonkin *et al.*, 2008) that suggests each of us will find (or the system will recommend) taggers whose judgement we feel we can trust.

Spamming and gaming search

Can a model succeed in which semantic annotations are automatically considered trustworthy? These differing methods pit authority against managed chaos. Given present text-based search technology, we do not face a lack of information, though we may lack trusted sources. Search engines can be fooled. A lucrative industry exists around the process of malicious gaming of search results. We face a world in which almost anything, anywhere, can be found, a vision which Morville described as ‘ambient findability’. Most things are ‘findable’. In the near future, the problem of search may become subordinate to the problem of filtering information, applying technologies initially designed, in the language of ubiquitous computing, to *augment* reality, to enable voluntary perceptual blindness – allowing us to direct our attention effectively to the information that matters to us.

How do we search?

Software agents are most capable when data placed before them is formally encoded in terms of semantics that they can relate to their own set. They do not, of course, ‘understand’, but depend on formal reasoning and the set of symbols available to them.

An agent might search as follows:

‘Find: [Any document] [about] [unique topic ID]’

receiving the response:

[Document at <http://www.uniquetopic.com>] [about][unique topic ID]...

The agent has received a readable response, is satisfied and processes these documents.

Is this an accurate model of user behaviour? Grosky (2002) notes that management systems for multimedia document collections and their users are typically at cross-purpose. Systems normally retrieve multimedia documents based on low-level features; users apply the abstract notion of *relevance*. Though classical, this is by no means the ideal model of user behaviour during interaction with a search engine.

Instead, successful search methods and interfaces are often designed to support specific aspects of information seeking on the Web. We seek known pages or other objects frequently (Choo, 2000), but Rieh (2004) found that for most people, we turn to the search engine when we cannot think of an appropriate topic-specific site. Queries *evolve* as the search progresses, terms are refined, and perhaps various different search engines with different capabilities are tried – for example, Rieh found that the ask.com natural language querying abilities are commonly used for some types of query. Ask.com is able to answer questions such as ‘How cold is Pluto?’ using a feature called ‘Smart Answers’ that provides answers at the top of the page. A version called ‘Ask for Kids’ [<http://www.askforkids.com>] provides a user-friendly question-refinement system enabling young people to refine their search (‘Tell me about the President’. Which one? ‘George Washington’). Wikipedia and its Simple English counterpart [<http://simple.wikipedia.org>] provide an analogous service.

Bates describes this model of search and retrieval as ‘berrypicking’; the user tries a succession of search terms, retaining useful or interesting results. By contrast, in ‘serendipitous browsing’, information is often found by happy accident, through semi-directed search or wider reading around a topic. Precision and serendipity do not go hand in hand. The most authoritative piece of information about a topic may not be the best resource. What about an opinion piece, a rant or a reaction, a thoughtful review or a beginner’s introduction offering a new perspective? In multi-disciplinary academia, establishing links between groups is difficult to do without going back to the basics, questioning and re-evaluating what we think we know.

The optimal user, like the optimal student or researcher, is information literate, patient and able, ready to engage in an interactive process of negotiation – and he or she is rare. A recent CIBER briefing paper (2008) suggests that today’s digital library users have very heterogeneous approaches to information seeking. They ‘skim’, viewing a few pages from an academic site, then leaving. They spend little time on websites, but ‘power browse’, looking for ‘quick wins’, a finding supported by studies such as Nielsen (2006), which showed that users browse pages by scanning them unevenly, focusing mainly on the top left of the page. The CIBER study suggests that the ‘Google generation’, who look to the internet as a primary resource for information, have a low appreciation of what the internet is, little ability to determine the trustworthiness or relevance of resources, low information literacy and poor understanding of their information needs. The extent to which these findings can be generalised is unclear, but the

paper adds that there is ‘little evidence...that Google generation youngsters are fundamentally “different” to what went before’.

Lecturer Tara Brabazon (author of *The University of Google*) has banned her students from using Wikipedia and Google in their first year of study, stating that these services offer quick answers but that dynamic and critical thinking skills are required before they are suitable for students’ use. First-year students are instead provided extracts from peer-reviewed printed texts.

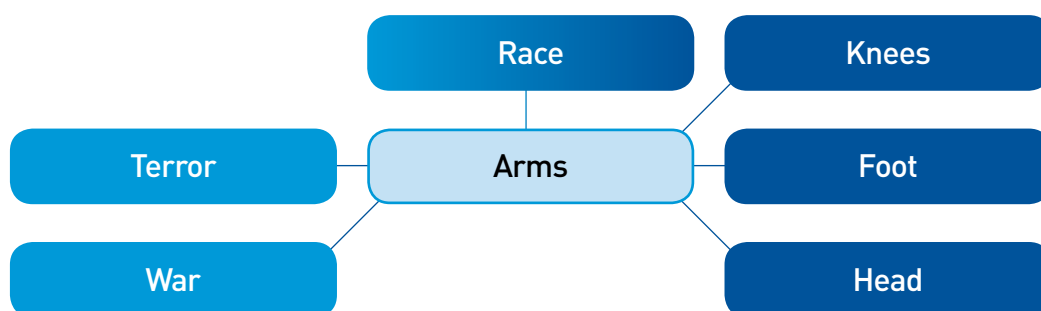
Data mining and artificial intelligence

The present state-of-the-art in artificial intelligence is not very good at understanding what things mean. However, computers are very good at brute-force calculation and abstract symbol manipulation. Anything possible in reasonable amounts of time by applying mathematical algorithms to a data structure is achievable by a computer.

Most of the statistical methods currently in use are based on some interpretation of the *distributional hypothesis*: that is, that two objects (such as words) are similar to the extent that they share contexts (Harris, 1968). That context might be adjacency within a sentence, co-occurrence within a sentence, paragraph, or document, or some other measure such as co-occurrence with another term.

For example, the words ‘for’ and ‘example’ frequently occur side by side. A document which frequently mentions ‘search’ might also contain frequent uses of ‘engine’, or possibly ‘rescue’, depending on the actual theme of the document.

Automated approaches can also distinguish between multiple uses of the same term. This is a major challenge for search engines such as Google: how does one distinguish between the use of the term ‘arms’, as in ‘arms race’, and its use in the sense of ‘arms, knees and head’? (Windows, 2004). This is actually one of the easier problems to solve via the distributional hypothesis; the use of the terms looks rather like this:



Where the word ‘arms’ co-occurs with ‘terror’, warfare may be involved. Where it co-occurs with ‘foot’, it probably refers to appendages. The role of statistics in guessing at meaning is in building useful generalisations based upon the available evidence. Given initial choices from both clusters, the reaction of the user may

permit the search engine to tailor further results according to the definition that has proven to be of interest. If they are exclusively interested in the topic related to warfare, results could be biased towards that definition. A similar approach can enable *entity disambiguation*, the ability to guess whether a particular use of the term 'George Bush' refers to the father or son. Co-occurrence is valid for words, but also for information like Web addresses and references to images or video.

There are many different approaches to statistical analysis. The mature technologies in use are not closely linked to human intelligence or judgement, though they may be inspired by research in those areas. The spam filter is probably the best-known example of supervised learning, based on a statistical trick known as Bayesian classification, described in Figure 3. Bayes does not return certainties, but a measurement of probability that a given object can be classified under a given label, given previous experience in the area.

Commercial search engines do not generally learn in the manner of a Bayes classifier, but work instead from ranking methods beyond the scope of this article. However, they too answer according to calculated probabilities: 'this query is similar to that article'.²

Figure 3: Bayesian reasoning in classification

Bayesian reasoning is probably the most familiar form of automated classification. It is mainly seen in spam filters. It is based around the following observation (Bayes' Rule):

$$P(\text{Cause}|\text{Effect}) = \frac{P(\text{Effect}|\text{Cause})P(\text{Cause})}{P(\text{Effect})}$$

This cryptic equation can be read in English as: (the probability that something that quacks is a duck) equals (the probability that ducks quack multiplied by the likelihood of there being a duck) divided by (the likelihood of hearing something quack). If it quacks like a duck and we have reason to expect to see a duck – and we have ample experience of the ways in which ducks quack to draw upon when making our judgement – then we correctly classify it as being a member of the genus *Anas platyrhynchos*. The advantage of Bayes is that it provides us with a quick and easy way of allowing a computer to make similar judgements. Search engines use many approaches, but in general the philosophical aim is similar: to determine, as cheaply as possible in terms of memory and computational power, the characteristics of the various objects on the web, the links between them, and their relationship to natural-language queries.

Multimedia search

The distributional hypothesis of semantics is not limited to text. For example, when an individual browses the Web, he or she travels on a *browsing path* from one page to another. Grosky (2003) suggests that two adjacent pages on that path of clicked links, particularly the part of the page surrounding the link that takes you to the following page, are probably similar in topic. Pages with similar layout are often assumed to be similar in topic; text elements that have similar layout may have a similar function, etc.

With data mining, we extract relevant features from a digital object, such as the key signature or melody of a piece of music (Hartmann *et al.*, 2007), the prevailing opinions in customer reviews (Hu *et al.*, 2004), or warning signs in lung cancer diagnosis (Perner, 2002). We can look for patterns and similarities between pages, or cluster web pages or other documents according to feature, such as the appearance of a 'sky-blue' colour in images.

Specialist search engine technology is moving ahead. In audio search, Nayio [<http://www.nayio.com>], though as yet unimpressive, holds out the promise of eventually being able to identify tunes by humming the theme into your microphone. Various video search services provide methods of searching by keywords spoken throughout the film. Google Video (Beta) searches the closed captioning of various television channels. Speechbot, a project run by Hewlett-Packard between 1999 and 2005, indexed over 17,000 hours of radio programming via speech recognition, but the project was terminated with the closure of HP's Cambridge research lab. Blinkx (founded in 2003) searches video files using a similar method, offering services to MSN, Live.com, Lycos, Infospace and AOL. PBS runs several similar archives, also using speech-recognition technology, that has been applied to index several PBS programmes over the last several years, such as News Hour [<http://www.pbs.org/newshour/video>] and the Mathline series of mathematics educational videos [<http://www.pbs.org/teachers/mathline/lessonplans/aboutvid.shtml>].

TASI's review of image search engines [<http://www.tasi.ac.uk/resources/searchengines.html>] primarily describes image search methods that rely on the text surrounding the image for information about the image content. Google's web spider (googlebot) does not analyse images directly. By comparison, picsearch.com searches on information taken directly from the image such as dimensions, file size, file type and colour information, all features of the image itself rather than the image context within a document. This approach has also been taken in video indexing, such as for example by Clippingdale and Fujii (2003) and Lee (2005); hybrid or *multimodal* approaches that index by combination of video and audio information (for example: faces and voices) have also been proposed in the research area (see for example Albiol *et al.*, 2004). IBM's 2006 Marvel search service (for local installation) uses advanced content analysis techniques for labelling and analysing data, and

provides hybrid and automated metadata tagging. Recent UCSD research, a supervised content-based image analysis for recognition of image elements (Carneiro *et al.*, 2007), demonstrates the complexity of the problem. Recognition of simple objects is shown to be possible with a fair degree of success, but complex objects or concepts represent a tougher problem.

Knowledge in context

Physical context

Another facet to the use of artificial intelligence in semantic annotation is that of *context-sensitivity* in mobile devices – devices that can sense their surroundings. Context-sensitivity – brought into the forefront of ubiquitous computing discourse by Dey (2001) – is yet another ‘idea’ technology. The aspiration is a ‘what you see is what you need’ experience, providing services and information without explicit prompting by the user. It is promising, but challenging in practice.

The device-level view of its immediate surroundings is quite different to that of the user. Though it is possible for a device to apply statistical/machine-learning methods in order to learn classifications, or for the developer to hardcode a set of rules defining the characteristics of a given context, it proves difficult to recognise the *same* contexts and features that the device’s owner considers significant. Although a fascinating research field and a technology which is bound to come to the fore with increasing popularity of mobile phones and other wearable devices, context-sensitivity is not a silver bullet. At present, most practical successes in context-sensitivity relate to position, physical state (body monitoring, such as heart rate) and physical activity (walking, running, driving a car).

A camera-phone might tag an image with the location at which it was taken, allow one to browse emails according to the location at which they were written or received, or provide relevant information on a just-in-time basis. Familiar scenarios in the latter research domain include *mediascapes* – contextual availability of information, such as information about pictures in a museum. Identifiers such as Radio Frequency Identification (RFID) tags or barcodes can be used as *keys* to retrieve relevant information, as in landmark projects such as Hewlett-Packard’s Cooltown (HP, 2001) and its descendants, such as BBC Coast and the open Create-A-Scape software³. Mediascapes are typically predefined rather than being dynamically generated from unstructured data, though both are possible. Services such as Panoramio, which interfaces closely with Google Maps, and Flickr, which permits geotagging, can also be used to create a ‘mediascape 2.0’ via user-contributed metadata.



³ http://www.createascape.org.uk/create_a_mediascape/make_mediascape/start_software.html



Several organisations have examined the possibility of location-based image search. These hold the promise of identifying landmarks by taking a photograph and sending it to an appropriate search service, which compares the subject of the image to reference images in order to provide an identification; see for example Davies *et al.* (2005) who describe a mobile tour guide that makes use of a 'point and find' mode of interaction; taking a photograph of an object causes it to be identified (as in Nokia's recent 'Point and Find' prototype). The system then provides information about that object. In practice, subjects often preferred browse modes that provided information about nearby landmarks without requiring the subject to specifically search for each one in order to retrieve information – as with the 'berrypicking' model.

Mobile devices discourage long search sessions due to transmission costs and ergonomic factors, but offer a great deal of information about users' context and information needs, of which commercial interface designers are only now beginning to take advantage. Groups such as GeoVector are exploring commercial use of 'point and click', and recent developments such as Google's free interface software for Google Maps, available for many Nokia devices, are beginning to change this trend.

Social context

When information-retrieval experts talk about context, they seldom think of location or physical variables, but the contexts in which a document was written, in which searches take place, or in which a document is retrieved. The way in which we apply or read symbols is dependent on these variables; both our use of language and search model shift with context of use. Teachers looking for appropriate resources to use in a classroom setting expect different responses to the pupil who types the topic into a search engine in the hope of finding an easy introduction. Context defines method, and is one of the most useful, important yet consistently overloaded and challenging words within and beyond the information retrieval world.

Search engines may provide a service tailored to many aspects of the context of use. Autology, for example, uses information about the user's present task – monitoring a document as it is being written – to tailor search results to the user's personal profile and deliver them 'just-in-time'.

Adaption, personalisation and social search

Searching is hard because there are two degrees of separation between the real world, the way in which things and experiences are represented by our brains or by a computer, and the many ways in which we use language to talk about them. As humans, we have inherited a 'theory of mind', that is, an outstandingly accurate ability to guess at what goes on inside other people's heads. This ability lets us guess at the right way to encode information in

language. We are so good at it that mostly, we are almost completely unaware of it. Computers, on the other hand, are 'mind-blind'. They don't know who we are or where we come from – our socio-cultural background, dialect or perhaps even language. Social network analysis – learning about the structure or society that created the data – is one means to solve this problem.

Computers are limited, simple-minded and cannot apply human standards of interaction. But the user may explicitly inform the computer of his or her preferences by customising their experience. For example, software can be customised to our language preferences, prioritising English results over other hits relevant to the query (although if we are bilingual this is possibly not our intention).

The second solution is the use of subtler forms of personalisation: recommendations based on explicit or implicitly given feedback – collaborative filtering and user profiling. Well-known examples include Amazon's recommendation system. Collaborative filtering looks at similarity; you are an individual and nobody else is quite like you, but there is a group of people who are more similar than others. If you are identified by a system which has learnt something about your preferences and interests as similar to certain other people's, the system can start to test that judgement by showing you the preferences of similar users ('You might like to watch this DVD'). Your reaction to that, if any, helps to narrow down your profile further. Feedback may be implicit (you ignored an option) or explicit (you requested not to see this again).

These methods exploit emergent patterns in society. Though unpredictable as individuals, as society comprises a set of stable structures emergent from individuals, so does our social behaviour and language. We cannot explicitly program the rules that we follow every day, but can generalise across the structures in the aggregate dataset.

Discussion

There is no single 'killer app'. Search engines are blunt instruments designed around the technically feasible. No search engine solves every class of problem with equal facility – probabilistic approaches are by necessity tuned to the average use case. Edge cases (unusual problems) may or may not generate appropriate results.

According to the research, finding the right information is a process, not a single interaction. Search engines cannot answer questions that we have not clearly articulated in our own thoughts, but we can expect some level of support in gaining a better understanding of the topic area, and perhaps can aspire to devices that offer information that we did not yet know we needed. The languages of search and interface will diversify further from the familiar text boxes and submit button, and towards a closer integration with ourselves and the ways in which we choose to work and live.

Mobile and ubiquitous computing hold promise in radically novel and perhaps more intuitive interfaces, such as tangible interfaces that let us explore information about our environment by manipulating elements within that environment directly. Indicating an object might provide us with information about it; building a structure of smart objects might affect our online environment.

It is always dangerous to offer predictions about the future. Successful services are usually those that incrementally enhance our lives – *plus ça change, plus c'est la même chose*. Social networks, mobile phones and the increasing availability of digital media link us more closely together and decrease the importance of physical distance. For this author, the ideal technology is one that performs an analogous task with information, linking search and browse, seamlessly relating different media and resource types, offering new information and starting points for discovery, without annoying the user, applying context-awareness to increase relevance and appropriateness, and rewarding curiosity.

The development of such a system is closely linked to awareness of the hidden structure and habits that make up our day-to-day lives. So next time you hum a tune and wonder about its name, wish you knew the location of the nearest Italian restaurant, or ask yourself why the sky is blue or what that book was you read last month, make a note: you have successfully identified a use case for contextual search.

Conclusion

One message is essential to understanding technology and integrating it into our lives. There is no Google generation and no shortcut to understanding. *People are still people*. With every advance there is a rush of hope and hype; will this technology transcend our limitations? Will the internet democratise society? Can the Semantic Web sweep the complexities of epistemology aside? Will Google teach us information literacy? The answer has always been 'No'. That is still the educator's domain.

The Web offers a wide expanse in which to lose ourselves, but searching and browsing are familiar skills, with familiar pathways for acquisition: attention span, problem-solving, knowledge integration and theory development. As technologies become more familiar to us, we may hope that both teachers and students may spend less time formulating the question and more time making use of the many rich resources on the Web, but computers do not replace motivation and cannot save us the trouble of understanding things ourselves. Indeed, technologies tend to highlight our shortcomings.

Rumours of the death of Education 1.0 have been greatly overstated. Whilst the increasing availability of information of all sorts and the increasingly flexible means by which it may be accessed and filtered are advances that will be welcome to most, this is evolution – not a revolution.

Bibliography


- Alberto Albiol, Torres, L., and Delp, E. J. (2004). 'Face recognition: when audio comes to the rescue of video'. In *Proceedings of the International Workshop on Image Analysis for Multimedia Interactive Services*, Lisboa, Portugal, April 2004.
- Allsopp, John (2006). 'The Big Picture on Microformats'. Digital Web Magazine. http://www.digital-web.com/articles/the_big_picture_on_microformats
- Bates, Marcia. (1989). 'The Design of Browsing and Berrypicking Techniques for the Online Search Interface'. <http://www.gseis.ucla.edu/faculty/bates/berrypicking.html>
- Berners-Lee, T., Hendler, J., and Lassila, O. (2001). 'The Semantic Web'. *Scientific American* 284 (5): 34-43.
- Berry, Michael W. and Browne, Murray (2005). *Understanding Search Engines: Mathematical Modelling and Text Retrieval*, Second Edition. Siam, Philadelphia. Society of industrial and applied mathematics.
- Carneiro, G., Chan, A. B., Moreno, P.J. and Vasconcelos, P. (2007). 'Supervised Learning of Semantic Classes for Image Annotation and Retrieval'. *IEEE Trans. Pattern. Analysis and Machine Intelligence*, March 2007 (Vol. 29, No. 3) pp. 394-410.
- Choo, C. W., Detlor, B. and Turnbull, D. (2000). 'Information Seeking on the Web: An Integrated Model of Browsing and Searching'. *First Monday*, volume 5, number 2 (February 2000). [http://firstmonday.org/issues/issue5_2/choo/index.html]
- Clippingdale, Simon and Fujii, Mahito (2003). 'Face Recognition for Video Indexing: Randomization of Face Templates Improves Robustness to Facial Expression'. *Lecture Notes in Computer Science*, Volume 2849/2003. Springer Berlin / Heidelberg.
- Davies, N., Hesse, A., Dix, A. and Cheverst, K. (2005). 'Understanding the Role of Image Recognition in Mobile Tour Guides', in *Proceedings of the Mobile HCI 05*, Salzburg, Austria.
- Dey, Anind K. (2001). 'Understanding and Using Context', *Personal Ubiquitous Computing* 5(1), 4-7.
- W.I. Grosky, D.V. Sreenath, and F. Fotouhi (2003). 'Emergent Semantics from Users' Browsing Paths,' *Proceedings of the First International Conference on Intelligence and Security Informatics, Lecture Notes in Computer Science*, H. Chen *et al.* (Eds.), Volume 2665, Springer Publishing Company. Berlin, Germany, pp. 355-357.
- Grosky W. I., Sreenath D. V., and Fotouhi, F. (2002) 'Emergent Semantics and the Multimedia Semantic Web'. *ACM SIGMOD*, Volume 31, Issue 4. Special section on semantic web and data management. Pages: 54 – 58. ISSN:0163-5808.
- Guy, M. and Tonkin, E. (2005) 'Folksonomies: Tidying up tags'. *D-Lib* 12(1). <http://www.dlib.org/dlib/january06/guy/01guy.html>
- Hewlett-Packard, (2001). Cooltown archive. <http://www.hpl.hp.com/archive/cooltown>
- Hu and B. Liu. (2004). 'Mining and Summarizing Customer Reviews'. In *KDD*, pages 168-177, Seattle, WA.
- Johnston, Amy E. Boyle, (2007). 'Ray Bradbury: Fahrenheit 451 Misinterpreted. L.A.'s August Pulitzer honoree says it was never about censorship'. *LA Times*, Wednesday, May 30, 2007. <http://www.laweekly.com/news/news/ray-bradbury-fahrenheit-451-misinterpreted/16524>
- Kipp, M., (2006). 'Complimentary or discrete context in online indexing: a comparison of user, creator and intermediary keywords'. *Canadian Journal of Information and Library Science* 30(3). <http://dlist.sir.arizona.edu/1533>
- Lee, Juhnyoung and Goodwin, Richard (2005). 'The Semantic Webscape: A View of the Semantic Web'. *WWW 2005*. <http://www2005.org/cdrom/docs/p1154.pdf>
- Lee, Jae-Ho, (2005). 'Automatic video management system using face recognition and MPEG-7 visual descriptors'. *ETRI Journal* 27 (6), December 2005
- Marshall, Catherine C. (1997). 'The future of annotation in a digital (paper) world'. Presented at the 35th Annual GSLIS Clinic: Successes and Failures of Digital Libraries

- McCool, R. (2005). 'Rethinking the semantic Web. Part I'. *IEEE Internet Computing*, Volume 9, Issue 6, Nov.-Dec. 2005 Page(s): 88, 86 – 87 Digital Object Identifier 10.1109/MIC.2005.133.
- Nielsen, J. (2006). 'F-Shaped Pattern for Reading Web Content'. http://www.useit.com/alertbox/reading_pattern.html
- Paolillo, John C., Mercure S., and Wright, Elijah (2005). 'The Social Semantics of LiveJournal FOAF: Structure and Change from 2004 to 2005'. ISWC 2005.
- Perner, P. (2002) Image mining: issues, framework, a generic tool and its application to medical-image diagnosis'. *Engineering Applications of Artificial Intelligence*. Volume 15, Issue 2, April 2002, Pages 205-216.
- Hartmann, K., Büchner, D., Berndt, A., Nürnberger, A. and Lange, C. (2007). 'Interactive data mining & machine learning techniques for musicology'. CIM 07.
- Rieh, Soo Young (2004). 'On the Web at home: Information seeking and Web searching in the home environment'. *JASIST* 55(8): 743-753 (2004). http://www.si.umich.edu/rieh/papers/rieh_jasist2004.pdf
- Shirky, Clay (2003). 'Semantic Web, Syllogism and Worldview'. http://www.shirky.com/writings/semantic_syllogism.html
- Tonkin, Emma., Corrado, Edward M., Moulaison, Heather L., Kipp, Margaret, E. I., Resmini, Andrea., Pfeiffer, Heather D. and Zhang, Qiping (2008). 'Collaborative and Social Tagging Networks'. *Ariadne* 54.
- Widdows, Dominic (2004). 'Geometry and Meaning'. CSLI Lecture Notes No. 172, Center for the Study of Language and Information, Stanford, California.

Interactive displays and next-generation interfaces

Michael Haller Upper Austria University of Applied Sciences

Until fairly recently, the limitations of display and interface technologies have restricted the potential for human interaction and collaboration with computers. For example, desktop computer style interfaces have not translated well to mobile devices and static display technologies tend to leave the user one step removed from interacting with content. However, the emergence of interactive whiteboards has pointed to new possibilities for using display technology for interaction and collaboration. A range of emerging technologies and applications could enable more natural and human centred interfaces so that interacting with computers and content becomes more intuitive. This will be important as computing moves from the desktop to be embedded in objects, devices and locations around us and as our 'desktop' and data are no longer device dependent but follow us across multiple platforms and locations.



The impact of Apple's iPhone and an increasing number of videos of multi-touch surfaces available on YouTube [http://www.youtube.com/results?search_query=multi+touch], show that users' expectations about using these devices in their daily lives have increased. The reaction to these natural interface implementations has been very dramatic. This is because people are still interested in a simpler way of navigating information and content where the computer interface is not a barrier, but enables them to accomplish tasks more quickly and easily. Multiple metaphors and interaction paradigms using pen, touch, and visual recognition are coming together with the other elements to create a new experience. In education, intuitive interfaces lower the barriers to using IT, allow for a better understanding of complex content and enhance opportunities for collaboration. In the near future it is likely that emerging display technologies such as electronic paper and OLED (Organic Light-Emitting Diode) screens will be delivered on flexible substrates. This will enable bendable/rollable displays that can be made larger than the dimension of the mobile device they are used with. E-paper could also enable inexpensive, very large digital displays to be incorporated into walls and other surfaces more widely. Speech recognition, gesture recognition, haptics, machine vision and even brain control are all improving rapidly to support more natural interactions with these new display technologies. This article concentrates on developments in different multi-touch surfaces and related applications. It also describes particular challenges and solutions for the design of tabletop and interactive wall environments and presents possible solutions for classrooms.

With the increasing development of interactive walls, interactive tables, and multi-touch devices, both companies and academics are evaluating their potential for wider use. Bill Gates noted at the 2008 Consumer Electronics Show that display technology is not just improving in quality, but also in the way that we interact with large surfaces. These newly emerging form factors require novel human–computer interaction techniques. Although movies such as *Minority Report* and *The Island* popularised the idea of futuristic, off-the-desktop gesture-based human–computer interaction and direct manipulation-based interfaces, in reality, making these interfaces is still a challenge. Conventional metaphors and underlying interface infrastructures for single-user desktop systems have been traditionally geared towards single mouse and keyboard-based WIMP (Window, Icon, Menu, Pointing) interface design. However, a table/wall setting provides a large interactive visual surface for groups to interact together. It encourages collaboration and coordination, as well as simultaneous and parallel problem solving among multiple users and therefore needs new kinds of interface.

Interactive surfaces

In late 1988, Xerox PARC developed the Live-Board¹, the first blackboard-sized touch-sensitive screen capable of displaying an image. Many in education will now be familiar with the interactive whiteboard. SMART Technologies Inc. [<http://www.smarttech.com>] introduced its first interactive whiteboard SMARTBoard in 1991. The tracking is based on the DVIT (Digital Vision Touch) technology and uses small cameras mounted in each of the four corners of the panel to track the user input². The system is mainly designed to be used with pens, but it can also track finger touches. A great number of digital whiteboards have also been sold to universities and educational institutions.

A similar technology is the touch frame provided by NextWindow. Again, embedded cameras track up to two points at the same time. The MIMIO and eBeam ultrasonic tracking devices, where participants use special styli, are a good and cheap alternative tracking surface. However, they are limited in their range, and line-of-sight restrictions reduce the tracking performance.

More recently, touch interfaces have been able to respond to multiple touches and gestures, increasing the possibilities for interaction and for multiple users to collaborate. Interactive tables, for example, have begun to move from prototype to product and combine the benefits of a traditional table with all the



¹ Elrod, S., Bruce, R., Gold, R., Goldberg, D., Halasz, F., Janssen, W., Lee, D., Mc-Call, K., Pedersen, E., Pier, F., Tang, J., and Welch, B., *Liveboard: a large interactive display supporting group meetings, presentations, and remote collaboration*, CHI '92 (New York, NY, USA), ACM Press, 1992, pp. 599–607.

² Morrison, G., 'A Camera-Based Input Device for Large Interactive Displays', *IEEE Computer Graphics and Applications*, vol. 25, no. 4, pp. 52–57, Jul/Aug, 2005.

functionalities of a digital computer. Although interactive tabletop environments are becoming increasingly common (see for example Mitsubishi Electric Research Lab's DiamondTouch, Microsoft's Surface and Philips' Entertable), there are few applications which fully show their potential. One area where they could be expected to be very useful is in supporting creative collaboration. In the creative process, people often sketch their ideas on large tables. A digital tabletop set-up could therefore provide an ideal interface for supporting computer-based collaboration. To better understand the design requirements for interactive displays in a business setting, we carried out an exploratory field study at Voestalpine, an Austrian steel company, which wants to use a tabletop surface for brainstorming sessions. We found the following design recommendations for an interactive, large vertical/horizontal display:

- Multi-point interaction and identification
- Robust tracking under non-optimal conditions
- Hardware robustness
- Physical objects should not interfere
- User can interact directly with the system
- Reasonable latency
- Inexpensive to manufacture.

We also noticed in our study that a direct touch on the surface seems to be an intuitive interaction metaphor. Especially at the beginning, users are not interested in using additional devices such as a stylus. On the other hand, the user's finger often obscures parts of the screen. Moreover, the screen gets dirty from finger prints. Albert noted that finger-operated touch screens are best for speed and worst for accuracy³.

DiamondTouch

Up to four users can sit on special chairs around the DiamondTouch⁴ table interface developed at the Mitsubishi Electric Research Lab (MERL). The sensing technology behind DiamondTouch is an XY pair of antenna arrays embedded in the surface of the table. Each user sits in a wired chair that broadcasts a unique radio signal. These signals are capacitively coupled through the user's body and into the antenna array whenever touches occur (Figure 1). Because each user sits in a different chair, the table is able to distinguish touches among the users.



³ Albert, A. E. The effect of graphic input devices on performance in a cursor positioning task. *Proceedings of the Human Factors Society 26th Annual Meeting*, Santa Monica, CA: Human Factors Society, 1982, pp. 54-58.

⁴ Dietz, P.H., Leigh, D.L., *DiamondTouch: A Multi-User Touch Technology*, ACM Symposium on User Interface Software and Technology (UIST), ISBN: 1-58113-438-X, pp. 219-226, November 2001.



Figure 1:
Two players sitting around
the DiamondTouch
playing 'memory'

The DiamondTouch is not only able to track multiple touches, but also able to *identify* different users (we can therefore call the system a multi-*person* system). The digital content is always projected onto the table's surface (see Figure 1). Another advantage of this table is the fact that additional objects placed on the surface do not interfere with the system. The interpolated resolution of the DiamondTouch is 2736x 2048 points (with a physical screen size of 42 inches) and the table can read out tracking information with a refresh rate of 30Hz. A similar set-up is presented by Rekimoto⁵ with the SmartSkin project, where he uses a mesh-shaped sensor grid to determine the hand position.

Microsoft's Surface

More recently, Microsoft presented the Surface table. This is expected to come to market later in 2008, but its price will initially limit its wider appeal. The system enables interaction with digital content through natural gestures, touches and physical objects. The Surface can track up to 40 simultaneous touches. In contrast to the DiamondTouch, the Surface is based on an optical tracking set-up, where five embedded infra-red cameras track the entire table (the current prototypes have a screen size of 30 inches). A special rear-projection surface and an embedded projector allow an optimal image. With the special projector, the engineers developed a relative low-sized table with a maximum height of 56 cm. The Microsoft team demonstrates the table's advantages with effective demonstrations developed for Sheraton Hotels, Harrah's Casinos, and T-Mobile. In the photo-sharing application, for instance, friends can put their WiFi digital camera on the table and share their photos in a very natural way (see Figure 2).

An alternative is to recognise and pair a device with RFID (Radio-Frequency Identification) tags or NFC (Near Field Communication). In this case, the table includes RFID readers which in combination with RFID tagged objects can be used to save and load different content. NFC allows devices to set up a link when brought together in close proximity. It is primarily designed to be used on mobile phones. The content, however, has still to be sent over Bluetooth (or another suitable link), since the NFC technology is not designed to transfer large amounts of data. RFID/NFC is likely to be included in increasing numbers of mobile phones and other devices, so in the future it may be possible for a user to have content from a mobile device appear on a large screen just by bringing their device within close range of the display.



Figure 2: All pictures are sent to the table's surface, once the Wi-Fi-based camera is put on the table. Alternatively, RFID tags on the devices can help for tracking devices on an interactive large surface.



Other interactive tables

Similar to the Microsoft Surface, the LumiSight table captures the objects on the table using cameras⁶ and a projector mounted inside the table. The InteracTable, a single-user system, allows interaction using a stylus. In contrast to related research, this system is based on a plasma display⁷. The DVIT cameras mounted in each of the four corners of the table track the users' input⁸. The lens of each camera has an approximately 90-degree field of view. The current version allows two simultaneous touches. Similar to the Microsoft Surface, people cannot place any physical objects (a coffee mug, for example) on the surface without achieving un-wanted touches. Stanford's iRoom table, an interface mainly designed for brainstorming discussion in schools, is another example, which is also based on the DVIT tracking with multiple DVIT frames.

One of the first larger tabletop set-ups has been presented by Ishii *et al.*⁹. In their installation, they implemented a set-up for engineers discussing urban planning. The system supports multi-layering of 2D sketches, drawings and maps in combination with 3D physical (tangible) objects, and is primarily designed for group sizes up to 10 people. The set-up consists of two projectors hanging from the ceiling. Two cameras (also mounted above the set-up)

⁶ Kakehi, Y., Iida, M., Naemura, T., Shirai, Y., Matsushita, M., Ohguro, T., 'Lumisight Table: Interactive View-Dependent Tabletop Display Surrounded by Multiple Users', In *IEEE Computer Graphics and Applications*, vol. 25, no.1, pp 48 – 53, 2005.

⁷ Streitz, N., Prante, P., Röcker, C., van Alphen, D., Magerkurth, C., Stenzel, R., 'Ambient Displays and Mobile Devices for the Creation of Social Architectural Spaces: Supporting informal communication and social awareness in organizations' in *Public and Situated Displays: Social and Interactional Aspects of Shared Display Technologies*, Kluwer Publishers, 2003. pp. 387-409.

⁸ Morrison, G., 'A Camera-Based Input Device for Large Interactive Displays', *IEEE Computer Graphics and Applications*, vol. 25, no. 4, pp. 52-57, Jul/Aug, 2005.

⁹ Ishii, H., Underkoffler, J., Chak, D., Piper, B., Ben-Joseph, E., Yeung, L. and Zahra, K., *Augmented Urban Planning Workbench: Overlaying Drawings, Physical Models and Digital Simulation*. IEEE and ACM International Symposium on Mixed and Augmented Reality ACM Press, Darmstadt, Germany.

capture all the users' activities. Finally, Han¹⁰ demonstrated an impressive scalable multi-touch interaction surface that takes advantage of frustrated total internal reflection (FTIR), a technique used in biometric applications such as fingerprint scanning. When light encounters the interface to a medium with a lower index of refraction, the light becomes refracted and beyond a certain angle, it undergoes total internal reflection. In contrast, another object (such as a finger) at the interface can frustrate this total internal reflection, causing a visible blob on the backface of the surface. This tracking system is highly scalable and very accurate – even under different lighting conditions.

As seen in this section, many companies and research laboratories are working on interactive tables, since they combine the advantages of a traditional table (face-to-face communication) with the advantages of a computer (easy archiving of data, and sharing of content for example). For education, both students and teacher could benefit from these devices, because they allow a very intuitive interaction with complex information.

Interactive paper

Many users still prefer real print-outs and paper to capture rough ideas. Instead of replacing the current environment, we propose an approach where we integrate traditional paper into the digital environment. The support of information exchange between computer and non-computer devices seems to be more and more important. In this context, the design of solutions that seamlessly bridge the gap between these two worlds is the key factor for practical applications.



Figure 3: Instead of using a tablet PC during a flight (a), users still prefer pencil and paper (b). In the 'Office of Tomorrow' project, users can go to the meeting and present their ideas to the audience either by transferring the real ink data to the digital whiteboard or by transferring printed information of the print-out to the digital whiteboard (c).

Sketching ideas and taking notes are basic tasks that are performed frequently in the phase of preparing or during a meeting or presentation. For this reason,

tablet PCs have been used as a good alternative to notebooks, because they allow an easy-to-use interface for sketching ideas. However, they are still too heavy and too big to be used in different environments (users still don't like to use a tablet PC during a flight for making a quick note, for example). This is the reason why paper still has a lot of advantages: it is light-weight, easy to navigate, people get a fast overview, it is easy to annotate, it is socially well accepted, and it doesn't require any power. The use of real paper and digital information combines the advantages of paper and additionally enhances them through the possibilities of the digital world. AceCad's Digimemo device [<http://www.acecad.com.tw/dma502.html>] for example, captures digitally and stores everything users sketch with ink on ordinary paper without the use of computer and special paper. Users can also easily view, edit, organise and share their handwritten notes on the PC.

More recently, Hull *et al.* presented 'Paper-Based Augmented Reality', an interactive paper. Users can simply get additional information (a website, for example) on a mobile phone by pointing the device at a printed website link. The advantage of their system is that they do not use real-time OCR (Optical Character Recognition) for capturing the text; instead they are matching bounding boxes of the blurred text captured by the mobile phone with the bounding boxes stored in a huge database. On mobile devices, characters are often so blurry that OCR is impossible. Identifying the bounding boxes around words is possible since the spaces between words and lines can still be distinguished. Their algorithm uses the number of characters in horizontally adjacent words as a feature.¹¹

Digital pens

In the project 'Office of Tomorrow', developed by the Media Interaction Lab [<http://www.mi-lab.org>], the authors combine traditional input devices such as pen and paper with the digital environment. To capture the ink on the real print-out, they use the Anoto pen [<http://www.anoto.com>]. Anoto-based digital pens are ballpoint-pens with an embedded (IR) infra-red camera that tracks the pen's movements simultaneously. The pen has to be used on a specially printed paper with a pattern of small dots with a nominal spacing of 0.3mm. Once the user touches the pattern with the pen, the camera tracks the underlying pattern. It can then derive its absolute coordinates on the pattern and send them to a computer over Bluetooth at a rate of 70Hz. Anoto pens with Bluetooth are available from Nokia (SU-1B), Logitech (io-2), and Hitachi Maxell (PenIT). From the pen, they receive the pen ID, the ID of the pattern sheet (each page has a unique pattern), and the position of the pen tip on the paper.



¹¹ Hull, J., Erol, B., Graham, J., Ke, Q., Kishi, H., Moraleda, J., Olst, D., Paper-Based Augmented Reality. In Proceedings of the 17th International Conference on Artificial Reality and Telexistence (Esbjerg, Denmark, November 28-30, 2007). ICAT '07. IEEE, 205-209.

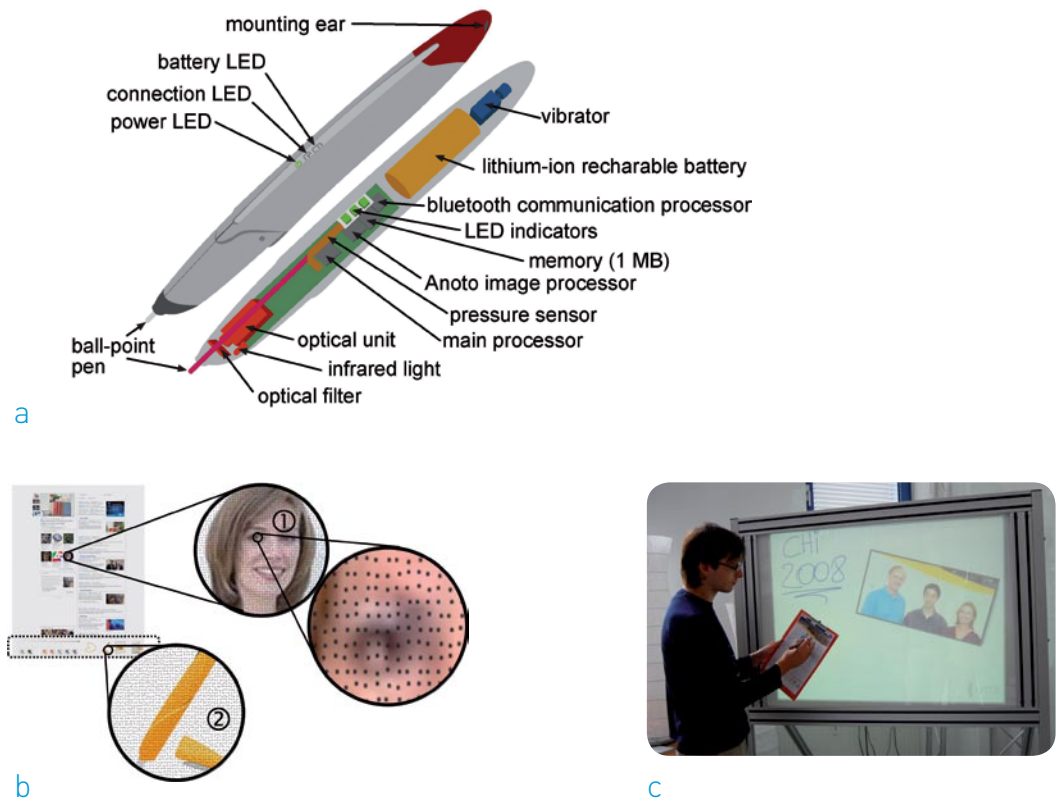


Figure 4: (a) The Anoto digital pen. (b) After exporting to an XPS file, the authors add an additional layer with two patterns on top of each print-out for tracking the strokes with the digital pen. While the upper part of the layer (1) is used for tracking the ink strokes on the page, the lower one (2) contains a unique ID for the control elements. (c) Elements from the real print-out can be sent to the digital whiteboard.

Each page that needs to become trackable has to be printed in combination with the pattern (see Figure 4b). The pen can be either used as an inking or pointing device that allows selections on the paper document and data manipulations on the digital whiteboard (Figure 4c). To change the mode for the pen, we integrated special control elements at the bottom of each page. By clicking on them, the pen can change its mode and selected data can be sent from the real print-out to the digital whiteboard. In addition, the system offers different options for defining the ink style including colours and stroke widths. By changing the colour or stroke width, only the digital ink will be changed accordingly, as the real ink still has the same colour or width. Our system supports additional annotations on the real print-out that can be performed with the real ink of the pen. The digital version of the ink can be either visualised in real time on the digital whiteboard or stored on the pen's integrated memory – the pen can store up to 70 written A4-pages. Real-time streaming is mainly used in scenarios, where the real print-out and the digital whiteboard are in the same location. Annotations on the paper are also immediately visible on the digital whiteboard. Again, the data transfer is accomplished through Bluetooth streaming from the Anoto device to the whiteboard PC.




Figure 5: (a) Annotations on the real print-out are immediately visible on the digital whiteboard. (b) Further manipulation of the data can be performed with the same digital pen.

Figure 5 depicts an example where a user is annotating with real ink on the paper document. The results are also visible as digital ink on the whiteboard and the audience can see immediately all changes done on the paper by the writer. All notes that are sent to the whiteboard can further be modified with digital ink (see Figure 5b). Transferred content such as images can be arranged and transformed on the digital surface. The sent data still has the same high quality as the item from the print-out (that is, the image from a website printed on the paper and sent to the digital whiteboard still has the same quality as the original image of the website).

Digital pens for learners

The Fly pen, manufactured by LeapFrog Enterprises, is an all-in-one solution for the digital Anoto pen designed for young children. There are a number of educational applications available for the pen, including a translator, a calculator and a writing assistant. Again, the pen uses paper with a pattern to track where the user writes on the page. A special object-recognition technology recognises what has been written and reads aloud the written symbols, characters and numbers. So, for instance, children can just sketch quickly a calculator on the paper and hear step-by-step advice on their math problems. Another very interesting pen has been proposed by LiveScribe whose Pulse device has an embedded OLED display and a microphone which can not only capture the handwriting, but also all users' audio comments. One model comes with 1GB and room for more than 100 hours of recorded audio, more than 16,000 pages of digital notes. Alternatively, a 2 GB model doubles the storage capacity.

Digital pens used at the interactive table and on a digital whiteboard

An alternative to using computer vision or other technology for hand tracking is capturing input through digital pens. The Shared Design Space¹², a collaborative interactive table, was the first demonstration that combined digital pens with surface tracking for a large tabletop set-up. Figure 6a shows our current prototype and illustrates the hardware setup of an interactive table and a digital whiteboard in combination with a rear-projection screen.

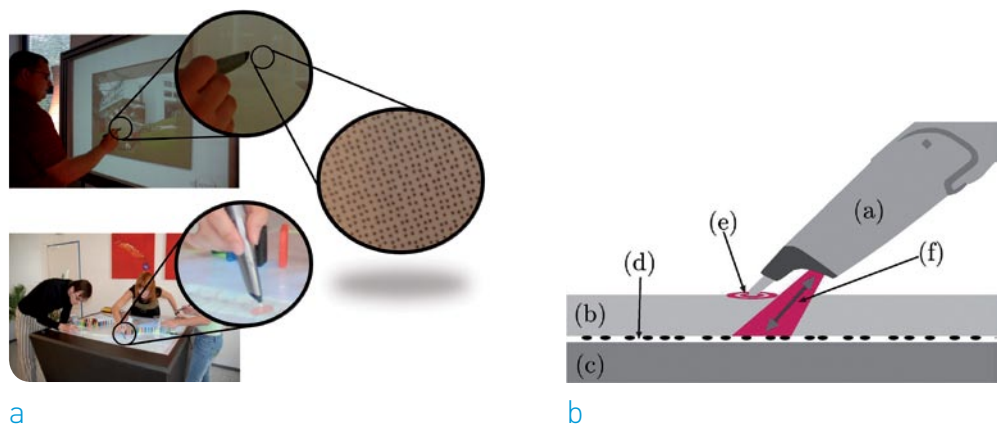


Figure 6: (a) The rear-projection screen has tiny dots printed on a special foil. (b) The different layers of the trackable table.

Figure 6b depicts the different layers used in our set-up. The digital pen (a) tracks the pattern, printed on a special backlit foil (d), which generates a diffuse light. Thus, no spotlights from the projectors are visible at the front of the screen. Moreover, the rendering and the brightness of the projected image are still of high quality. The set-up used one A0 sized pattern sheet (118.0cm × 84.1cm). The pattern is printed with the black ink cartridge (which is not IR transparent and therefore visible for the IR camera). Notice that the colours Cyan, Magenta, and Yellow (even composed) are invisible for the IR camera. The pattern is clamped in between two acrylic panels (b) (c). The panel in the back has a width of 6mm and guarantees a stable and robust surface while the panel in the front has a width of only 0.8mm to protect the pattern from scratches. The acrylic cover in the front does not diffract the pattern at all. However, using thicker front panels (≥ 4 mm), produces bad tracking results.

Digital whiteboard

Similar to the rear-projection table, the system has also been implemented successfully as a digital whiteboard using the same surface layers. In addition

to the pen tracking, our system also supports hand tracking. Behind the display surface, a common PAL camera has been mounted and tracked the user's hands on the screen by using brightness differences. The camera behind the screen captures a grey surface and objects coming near the surface will appear as blurred shadow. Thus, only objects directly touching the surface are recognised as sharp outlined shapes.

Conclusions

Multi-touch and interactive surfaces are becoming more interesting, because they allow a natural and intuitive interaction with the computer system.

These more intuitive and natural interfaces could help students to be more actively involved in working together with content and could also help improve whole-class teaching activities. As these technologies develop, the barrier of having to learn and work with traditional computer interfaces may diminish.

It is still unclear how fast these interfaces will become part of our daily life and how long it will take for them to be used in every classroom. However, we strongly believe that the more intuitive the interface is, the faster it will be accepted and used. There is a huge potential in these devices, because they allow us to use digital technologies in a more human way. We are just at the beginning of a new decade, where books can be displayed on e-paper devices such as the Sony Reader.

On the other hand, we will still work with traditional interfaces including paper. The integration of real notes, for example, in a digital environment seems to be a very important motivation for people using these new technologies, since it combines the affordances of a traditional medium such as paper with the capabilities of digital content and displays.

© Copyright Becta 2008

You may reproduce this material, free of charge, in any format or medium without specific permission, provided you are not reproducing it for financial or material gain. You must reproduce the material accurately and not use it in a misleading context. If you are republishing the material or issuing it to others, you must acknowledge its source, copyright status and date of publication. While great care has been taken to ensure that the information in this publication is accurate at the time of publication, we accept no responsibility for any errors or omissions. Where a specific product is referred to in this publication, no recommendation or endorsement of that product by Becta is intended, nor should it be inferred.

Millburn Hill Road
Science Park
Coventry CV4 7JJ

Tel: 024 7641 6994
Fax: 024 7641 1418

Research email: emtech@becta.org.uk
Main email: becta@becta.org.uk
URL: www.becta.org.uk

