

Research Report

March 2008

**Becta** leading  
next generation  
learning

# Emerging technologies for learning

## Volume 3 (2008)




# 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.<sup>1</sup> 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.)



<sup>1</sup> 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.<sup>2</sup>

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.



<sup>2</sup> 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 deskillng 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.

© 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](mailto:emtech@becta.org.uk)  
Main email: [becta@becta.org.uk](mailto:becta@becta.org.uk)  
URL: [www.becta.org.uk](http://www.becta.org.uk)

