Robots, drugs, reality and education: How the future will change how we think

Richard Sandford

About the author

Richard Sandford is a Learning Researcher at Futurelab, a not-for-profit organisation investigating the relationships between education and digital technologies. He led the research-gathering and scenario construction work that took place within the DCSF-sponsored ‘Beyond Current Horizons’ project, a foresight programme considering possible socio-technical futures to 2030 and beyond and their implications for education. Prior to that his research has concentrated on mobile technologies and the potential of games to support learning, working with government and commercial partners in Singapore and the UK.

http://www.beyondcurrenthorizons.org.uk/
http://www.futurelab.org.uk/

Contents

Introduction

Anybody who works in education is used to hearing about the future. A central purpose of formal education is to prepare learners for their future lives, to attempt to give them the skills and dispositions that will enable them to make a living and enjoy a fulfilling life. Perhaps because of this, ‘the future’ figures in most discussions of education, in a number of ways. It can be a destination to work towards; it might stand for a time when the problems of today will have been addressed; it might be thought of as a threat to react to, or as a source of inspiration, where the daily challenges and constraints of teaching seem remote, and new ideas can have room to breathe. It’s frequently – perhaps normally – associated with new technologies and an idea of ‘progress’ that’s been linked to engineering and scientific development since before the last century, and it’s been associated with certain ideas since even earlier times. More efficient – cleaner, faster – ways to travel will be developed. Entertainments will be more spectacular or life-like. Information will be more easily shared and accessed; remote communication with colleagues and family members will make us feel they are in the room with us. The things that we find difficult today will become easier.

From an educational perspective, these technological developments are always thought to presage change. New skills will have to be taught to students, in order to give them access to new employment markets created by these technologies. New pedagogic opportunities will be offered by new technologies for communications and display. We will have to do our very best to keep up (there is never any suggestion that education might be in the vanguard of these changes) and to make sure that we are not accused of being behind the times.

But is there a danger that we take our ideas about the future for granted? That we spend so much time hearing about the future that we never stop to check whether our assumptions match reality? Commonly accepted ideas about what’s coming next are powerful images that become more ingrained with each repetition: the ‘orthodox future’ (in a phrase coined by Peter Schwartz[1]) is hard to resist, and shapes our actions in the present often without our realising. It’s easy to forget, in the face of such consensus, that the future hasn’t happened yet, and when it arrives it might look quite different to the assumptions we made. There are plausible counter-possibilities for many of the trends that are often spoken about as if they are already facts. China may not become the new global power it’s often assumed to be if it is unable to find solutions for the erosion of productive agricultural land, or reach an effective political settlement between its coastal middle classes and inland border populations. Increased global access to the internet may reinforce national boundaries, making it easier for closed markets to function, rather than supporting a global conversation. Employers might find it more efficient to invest in the education

[1] Schwartz, P., *The Art of the Long View*, John Wiley & Sons, 1997

of their employees directly, rather than working with schools to ensure school leavers have appropriate skills for the workplace. For every uncertainty of this sort, there are different potential outcomes: there is thus not one single future but many possible futures. And our orthodox future is just one of them.

So perhaps it's worth taking a step back from the assumptions about education and the future that we're surrounded with, and re-examining some current areas of technological development that are likely to challenge some of our present-day beliefs about the nature and purpose of education. This article considers four broad areas of technological development over the next few decades, to 2030 or thereabouts, and for each one suggests what implications they might have for education. Some of these areas might seem well-rehearsed topics of discussion – but their familiarity doesn't prevent them developing in ways that we might find unexpected[2].

The technological directions of these various trends are very familiar, but their effect on the ways in which we talk about the world and the ways in which they shape aspects of our inner lives are less often considered. In each case, we'll consider the social changes that might shape its development – discussions of the future often seem to assume that technological change drives social change, but of course society influences technology as well, through setting research agendas, channelling research funds, inspiring invention through arts and culture, deciding which technologies are useful (it was never imagined that SMS messaging would find such a place in people's lives) and which are not (the market for immersive virtual-reality headsets turned out to be far smaller than manufacturers had calculated), and evolving moral and behavioural norms that govern its use. The impacts of new technologies lie not only in offering ways to do things faster or more easily, but also in influencing the ways we see the world. So the social context is important.

As with other assumptions about the future, these descriptions make a lot of general assumptions about the world: they all assume that increases in computing power continue, that mineral resources needed to support their development remain available, that the consumer economies necessary to support the development of these technologies are stable, and so on. And it is important also to note that they are not intended to be read as predictions. The intention is to raise questions about the ways in which technology and society interact, rather than making specific claims about what is or isn't likely. They are all drawn from recent research, and describe trends considered likely to play out over the next twenty or thirty years, as identified by informed commentators rather than drawn from the far reaches of science fiction.

[2] This article draws on the work of Professor Dave Cliff and colleagues, undertaken for the Beyond Current Horizons foresight programme [http://www.beyondcurrenthorizons.org.uk/]; their original article is Cliff *et al*. (2008), 'Future issues in socio-technical change for education' [http://www.beyondcurrenthorizons.org.uk/wp-content/uploads/socio_technical_change_revised_disclaimerv2.pdf]

It's worth remembering, of course, that mobile telephones' computing power or synthetic biology might have seemed just as far-fetched twenty years ago.

Artificial intelligence and robotics

Machines that have human characteristics have been a feature of legend and the goal of engineers throughout history, and perhaps the defining human characteristics are intelligence and/or consciousness. Certainly the notion of a single artificial entity able to think for itself, to reason and to display general intelligence, has been a stock motif for writers and film-makers for decades. However, there has been no sign of this kind of 'true' artificial intelligence despite over fifty years of research[3], and it is reckoned unlikely to appear within the foreseeable future[4]. That's not to say, however, that other forms of mechanical intelligence don't already have an impact on the way we live, or have the potential to change the way we interact with the world. The change of focus within AI research from creating general intelligence to solving specific tasks and problems, coupled with increased computing power, has led to a number of different applications across a wide number of domains, all of which can give the impression of possessing and using some kind of intelligence. The question facing us today is not whether we can create human-like intelligence, but what the risks are when humans mistakenly see intelligence where it doesn't exist.

At the most basic level, automated data-mining software based on principles of 'machine learning' is part of many people's web experiences, whether it's Google asking whether you meant to search for something else, Amazon letting you know what other people bought, or Flickr showing you the most 'interesting' photos in your search results. Presenting data in this way might seem a long way from thinking about robots, but it still raises issues of control, privacy and context. Having a personal data concierge to help you find your way through complex data sets (or just to fill out web forms for you) might be one thing, but having your data trail monitored automatically in the interests of a private corporation would be another. These software bots are used by many companies to automatically assess credit scores, manage call-centre queues (and, of course, to send spam). If these lines of code are taking decisions that affect human lives in some way, would it be more appropriate to have a human being involved? Or will we become more used to the idea, and realise that there might not be any loss of agency involved in handing such routine decisions to a software agent? In either case, it seems worth noting that this agency doesn't reside in a single body but is distributed across a network: as humans, it seems easier for us to think of intelligence as somehow belonging to an individual entity.

---

[3] The academic field of artificial intelligence is generally reckoned to have begun in 1956 at the Dartmouth Conference (see also this proposal from John McCarthy, one of the founder of the field [http://www-formal.stanford.edu/jmc/history/dartmouth/dartmouth.html]).
[4] Cliff *et al*. (2008)

Of course, at the level that most of us encounter these bots (on eBay, on price comparison sites and so on), questions about their autonomy are rather abstract, given their limited capabilities. Questions about how far decisions affecting human lives should be handed to machines become more urgent when these machines are used in warfare. Robots have been used on the battlefield for clearing mines and reconnaissance for a number of years, but since 2001 ‘Unmanned Aerial Vehicles’ (UAVs) have been used to attack enemy combatants, piloted remotely from air bases that may be thousands of miles removed from the theatre of war. This has raised questions about the degree to which human operators are informed about the conditions surrounding the lethal decisions they make. Currently, research is under way to enable these UAVs to be autonomous, operating free of any direct human control: the questions of responsibility and morality raised by the possibility of these robots fighting on our behalf are magnitudes greater. Efforts to equip these robots with a ‘conscience’ would raise as many questions as they answer.

Not all appearances of autonomy bring with them such weighty questions or are so visible, though they might be on a similar scale. Many autonomous robots work in industry. Increasingly, however, examples of autonomous machine intelligence are being found in domestic settings: since 2002, the Roomba, an autonomous vacuum cleaner, has sold over two million units, while the introduction of the Tivo in 1999 opened the door to digital video recorders that take decisions about the kinds of programmes the household is likely to be interested in. Perhaps less visibly, characters in computer and console games are increasingly likely to use some form of programmed ‘artificial intelligence’ to ensure the players’ interaction with them is convincing and sufficiently challenging. These three different sorts of intelligence are already embedded within the domestic sphere, and may pave the way for the acceptance of robots that act as companions for the elderly or ill, making it possible for members of vulnerable groups to be left in the care of an autonomous machine without the risk of censure.

Sometimes it might be possible to perceive the cumulative action of many smaller intelligences as a single artificial intelligence. Safety systems within cars are becoming sufficiently advanced that they are increasingly given responsibility for handling critical events, detecting and responding to the proximity of obstacles many times faster than a human can achieve: the cumulative effect of many specialist systems – the anti-lock brake system, traction-control system, blind-spot detector, lane-departure warning, emergency brake assist, and so on – may be to give the impression of a single general intelligence. Related to this idea is the notion of ‘crowd-sourcing’: while it may seem strange to include human intelligence in a discussion of technological intelligence, the cumulative effect of many minds (for example, within Amazon's HIT[5] system) can give the impression of a single intelligence.

All these examples suggest that the categories of intelligence and the ways in which we understand agency that we inherit from our human history are unlikely to work well as signposts in the future. 'Artificial intelligence' as a recognisable entity is not likely to be a feature of our future: instead, a little sprinkling of AI will be shared around almost everything, with small intelligences monitoring, calibrating and otherwise tending to their patch of the network. We will get used to our energy grids, traffic management, media consumption, our clothes and our furniture being 'smart', and more used to outsourcing cognitive jobs that are better managed by computers.

Implications for education

The most visible appearance of this kind of outsourcing within education is the automated marking of exam scripts. Approaches are being trialled in which language-parsing software is trained by human judges to recognise correlations between certain linguistic constructions and high scores, allowing a more standardised approach and ironing out inconsistencies between markers who may be tired or unduly subjective. Using technology to speed the process of giving students feedback on their work (and lessening the load on overworked teachers and academics) seems like a positive development. But there are concerns that machine marking will not be sensitive to the nuances and individuality that mark good writing, and that students will learn how to write for computers rather than people. This is an area in which the relationship between types of intelligence is still being negotiated, with compromises being tested, such as the Open University's approach of marking first by machine, then by human. Other possibilities include systems that are able to adapt to the learner's input and offer appropriate feedback (perhaps drawing on artificial intelligence processes developed by the games industry). Whatever is eventually settled on, there is little chance of machine intelligence not being a feature of educators' working lives over the coming years.

More broadly, questions that are currently thought of as abstract and philosophical will take on a new life in the light of machine intelligence, and learners will need new ways of thinking about what it means for an entity to be conscious, what the nature of personhood is, how far we can expect an autonomous robot to take responsibility for its actions and how far we should entrust one with things that matter to us as humans. As a minimum, learners will need to understand that what seems like intelligence might not be: being aware of the confusion that can arise from becoming attached to a personal communications assistant, or treating a car as a person, will be an important skill. In particular, educators will need to remain aware of the

[5] Amazon offers the paid opportunity to complete 'Human Intelligence Tasks', or tasks that people are currently better at than computers, such as recognising images, through its Mechanical Turk scheme [http://aws.amazon.com/mturk/].

sources of data they use to form opinions of learners and their progress and to have considered the ethical ramifications of judging a child on the basis of a software agent's report.

Psychopharmacology

Drugs that affect cognitive behaviour have, outside the recreational exploration of narcotics and so on, traditionally been developed to address conditions such as Alzheimer's disease, narcolepsy or ADHD: that is, developed with a medical purpose in mind. These treatments frequently target the symptoms of such conditions, such as memory loss, lack of attention, anxiety or apathy, and increasingly people without these conditions are finding the cognitive effects of the treatments – improved memory, greater concentration span, clarity of focus – attractive. Anecdotal evidence[6] suggests that many people, particularly in cognitively demanding roles such as engineering or medicine, are taking these 'smart drugs' or 'cognitive enhancers' for cosmetic reasons.

There are currently many unanswered questions around the practicalities of cosmetic pharmacology – the degree of regulation of suppliers, the health risks involved in taking a drug intended for narcolepsy when you don't have narcolepsy – but the general trend towards using medical technology to augment existing cognitive capability is established, and raises some difficult questions, as well as offering some opportunities. Most obviously, in the context of an ageing demographic in a country still in recession, cognitive enhancements might make it possible for established members of the workforce to remove the competitive advantage of newer members (their greater mental agility) while retaining the advantages of greater experience and extending the length of time they are able to remain productive (whatever happens to the state pension age, many will prefer, or at least find it necessary, to remain in employment).

Of course, this sort of wholesale medication of society would give the pharmaceutical companies licensing 'smart drugs' a different sort of role in public life. More fundamentally, the notion that prosthetic interventions of this kind are a normal part of life challenges our historic understanding that prostheses are remedial – whether physical or pharmacological, they are intended to bring the recipient up to a 'normal' level, at which point the prosthetic is no longer necessary. For the first time, prosthetic technology, applied inside or outside the body, is capable of helping the recipient perform at levels beyond the normal. Some groups might welcome this: those outside this definition of 'normal' have often tried to highlight the widespread attitude that sees unusual biological or cognitive configurations as deviations (for

[6] For example, *Nature* survey, cited in Cliff *et al.* (2008), Johann Hari [http://www.johannhari.com/archive/article.php?id=1298] and the recent *Guardian* article, 'Students turn to 'smart drugs' to boost grades' [http://www.guardian.co.uk/education/2009/oct/01/students-smart-drugs-boost-grades].

example, the efforts of many autistic people to promote the word ‘neurotypical’ in preference to ‘normal’). When this fact is encountered in a competitive environment, however, issues of fairness and cheating arise. This is problematic enough in a sporting context such as the Olympics, when the individuals are already far from typical; when the competition is for employment, any suggestion of an unfair advantage will be more keenly felt. Economic incentives to use these drugs (and hence supply them) may lead to unsafe working practices; more generally, those who are biologically unsuited to using them may find that this constitutes a disability.

Implications for education

These debates are likely to be felt most immediately within the education sector in the context of the medicalisation of childhood, the current concerns about the readiness of some groups to label some children’s behaviour as non-normal, and the use of medication to control behaviour, with all the attendant ethical issues these bring. The difficulty of defining cheating or what constitutes an unfair advantage in such a context is already clear: the widespread acceptance of cosmetic cognitive enhancement would place these issues in an even starker context. It is possible to imagine an examination room in which two children are taking the same drug, yet one of these has been diagnosed with ADHD and is entitled to the help the drug affords them, while the other is at an advantage and consequently cheating. At the root of these debates will be a discussion of what makes someone human: for there to be a point at which the effects of a drug cease to be remedial and become instead an advantage, there must be an agreed ‘standard human’. For most of our history this has been defined by biology; over the next few decades we will have to define it ourselves and be prepared to regulate it if we want to keep the forms of assessment and recognition we currently have. Perhaps these drugs will be the catalyst for the new forms of assessment that are regularly demanded at education conferences.

Of course, this unfair advantage may be reproduced globally: if the UK chooses to ban the use of cognitive enhancement drugs, how will our engineers and scientists compete in a global marketplace where not every government feels the same way? Or conversely, if the UK’s knowledge workers reap the benefits of them, does that make a mockery of international efforts to address the disparity between rich and poor nations? If some countries are struggling to find more urgent medication, they stand little chance of accessing (genuine, regulated) cosmetic pharmaceuticals. Smart drugs have the potential to highlight inequalities on a global and a local scale: just as it seems plausible to imagine that developed nations’ workforces will make use of these drugs ahead of poorer nations, so it seems likely that they might become another advantage enjoyed by the middle class, along with private tuition and higher aspirations.

More prosaically, schools will have to prepare for a workforce in which these drugs are used: this means not only making sure learners are aware of whatever social practices evolve around the place of these drugs in the workforce, but educators will have to establish their response to the attractions of these drugs. However attractive increased energy and alertness might sound to a harassed staff room, heads might be more aware of the need to ensure that parents don't see a teaching workforce assisted by the products of the big pharmaceutical companies as a dereliction of their vocation.

Augmented reality[7]

The capacity of technology to allow us access to what feels like a different kind of reality, whether called 'cyberspace' or a 'virtual world', is well established in our ideas of technology. This idea is often accompanied by the notion of a computer acting as a kind of portal or gateway between the real world and the virtual. But since the early 1990s, technologists and researchers have been thinking about standing this idea on its head: what if the digital world was present in the same place as our real world? The increasing mobility of our computing devices, coupled with the widespread access we (usually) enjoy to wireless data networks, has eroded the image of the desktop computer and changed our expectations of access to information: we expect to be able to participate online in many different places, rather than being tethered to one location. Additionally, augmentation technologies such as radio-frequency identification (RFID) and two-dimensional barcodes enable data to be attached to physical objects, meaning that the real world can have a presence within the virtual world. Rather than choosing between spending time either in the 'real' or the 'virtual' world, then, the possibility arises of blending these two options, tying digital information to physical locations in the real world: augmenting our reality with data from the network.

Researchers have been exploring what this sort of ubiquitous or pervasive computing might look like for a number of years, and education researchers in particular have been exploring the ways in which mobile technologies might support learning. What is different about the present moment is that the components necessary to pinpoint a device in physical space – a GPS-enabled device, wireless network access, a compass, accelerometers – have progressed from the bulky and often unreliable versions available to early researchers to the smooth and (mostly)

---

[7] Adam Greenfield has written a detailed overview of location-based and context-aware computing and described some possible implications for education in his Emerging Technologies article, 'Location-based and context-aware education: prospects and perils' [http://partners.becta.org.uk/upload-dir/downloads/page_documents/research/emerging_technologies_08_chapter3.pdf]. Steve Benford has written an overview of location-based games and projects, some with an educational focus, in 'Location-based technologies for learning' [http://emergingtechnologies.becta.org.uk/upload-dir/downloads/page_documents/research/emerging_technologies/locationbased_technologies_learning.pdf].

dependable features recognisable to anyone with an iPhone or recent Android device. Rather than being found within specialised equipment in the hands of a few technologists, these location-aware capabilities are now embedded in general devices that are in the hands of millions across the globe. This infiltration into the everyday makes possible the evolution and negotiation of norms and rules of use across whole societies, establishing the social practices and affordances of location-aware information technology. Over the next few months and years, designers and developers will be able to respond to authentic patterns of use derived from real-world practice, gaining an idea of the ways in which these technologies have demonstrated their relevance to people's lives, and of the compromises and accommodations people have reached with some of the more problematic features of this new augmented reality.

These features will go beyond the practical or commercial, challenging some of the deeply-held assumptions that underpin our social interactions. In particular, the idea that information is the same for everyone everywhere will no longer be a useful assumption, as individual contexts are increasingly taken into account by designers and developers. Although truly context-aware computing is an intractably problematic proposition, due to the difficulty of modelling social context and interior states of mind in a computational way, personal histories of behaviour and consumption already inform many of our experiences on the web, and there is little reason to think that this will not also be the case within augmented reality applications. It will be entirely possible to imagine two individuals in the same location, using the same application on the same mobile device, experiencing completely different places as their devices present them with a unique informational patina knitted together from their friends' activities, their own previous activities in this space, their cultural and commercial preferences, and so on. This is not only an issue for software developers: architects will have to remember that their buildings will be perceived through this informational haze, as will anyone who aims to address more than one person at once through their work.

Implications for education

As well as the obvious and well-documented potential these mobile and location-aware devices have for supporting new ways of learning[8], educators will need to be aware that new information literacies will be needed to make sense of these new forms of sharing information and to take full advantage of the educational opportunities they offer. Some of these might be thought of as 'physical literacies',

[8] For example, see Naismith *et al.* (2004), *Literature review in mobile technologies and learning*, Futurelab, UK [http://www.futurelab.org.uk/resources/documents/lit_reviews/Mobile_Review.pdf], Ally, M. (ed.) (2009), *Mobile Learning: Transforming the Delivery of Education and Training*, Athabasca University [http://www.aupress.ca/index.php/books/120155] and MoLeNet [http://www.molenet.org.uk/].

given the embedded and embodied nature of these devices and the technologies to which they afford access: understanding how physical action and embodiment influence understanding will be a priority for researchers. This embodied nature sets these technologies apart and presents a brand-new set of challenges for educators, ones that will need input from experts in drama and physical education as well as technologists. For practitioners, understanding that the reality they assume their students are familiar with may not be shared will be crucial, as will negotiating the use of augmented reality within schools: education will have a role to play in helping people to explore 'un-augmented' experience and safeguarding the notion of a common, shared experience.

Computing as bioscience

Although the development of silicon chips for computing has been accurately described by Moore's Law[9] over the past few decades, with processing power roughly doubling every two years or so, at some point in the next ten or twenty years physical constraints will prevent engineers squeezing ever-increasing numbers of components onto a chip. As the scale of these components approaches the level of individual atoms, quantum tunnelling effects are expected to affect the ability of electrons to move predictably and reliably; at such scales, too, managing energy consumption to reduce heat effects while providing sufficient power becomes problematic. While developments in nanotechnology have the potential to extend the silicon age, the inevitable physical limits constraining chip technology, coupled with the pressure for devices to use energy more efficiently and the diminishing availability of the (toxic) raw materials needed for their manufacture, will encourage the development of alternative forms of computing. Two technological alternatives currently being explored are optical computing (in which light, rather than a stream of electrons, is used to carry digital information) and quantum computing (a largely theoretical approach in which the capacity of subatomic particles to exist in multiple states is used to enable parallel information processing). A third approach looks to the biological sciences for inspiration.

Computer scientists are drawing on biology in two main ways: looking for the ways in which natural phenomena can inspire new approaches to computational challenges, and using biological material to perform computational operations. Neurological processes have shown researchers new approaches to understanding networks, while studying the ways in which populations evolve and groups of animals flock together has led to the development of new ways of refining algorithms and optimising solutions. The efficiencies delivered through these advances hold out the possibility of working with extremely large numbers and addressing complex informational challenge; nevertheless, they are usually employed with silicon-based

[9] See http://www.intel.com/technology/mooreslaw/ for more information.

computers. A more radical influence from biology can be seen in the use of biological materials as the physical basis for computing.

There are many different approaches being taken in the exploration of biologically-grounded computing, across a number of different domains, and the speed with which the area is developing makes it impossible to offer an accurate snapshot of all of them. In general, though, approaches seem to fall into those groups using neurons from basic life forms, such as leeches or slime moulds, in neural networks, and those that rely on the interactions between molecular chains. Some approaches to molecular computing are based on the capacity of certain chemical reactions to produce wave phenomena, whose interactions can be treated as logical operations (if two waves meet and reinforce each other, that might be thought of as an AND gate, for example)[10] and whose outcomes (such as the presence or absence of a precipitate) can be taken to stand for particular values of variables. Or they might be based on the predictable and regular interactions between strands of DNA (a ready-made system for encoding information), representing the problem to be solved by structuring the DNA strands in particular ways. Other molecules have been used by researchers, but they follow the same principle: suspended in a non-reactive solution, the molecules collide with each other, combining to form new molecules which represent possible outcomes. These are then weeded out through a series of chemical reactions to provide a single molecule whose structure encodes the correct answer[11].

At present, this is a cumbersome approach: the preparation for this process occupies many hours of researchers' time, as does the final process of extracting the solution. The appeal, however, lies in the speed with which the various possible solutions are discovered: there are trillions of molecules all colliding with each other at the same time, meaning that the computations are carried out in parallel rather than in sequence and so providing a massive increase in speed. Biological approaches towards computation also offer a number of practical advantages over traditional silicon-based computing. These approaches are more energy-efficient, losing much less heat than conventional computing, and while for domestic computing this is not an obvious disadvantage, on a commercial scale, the data centres on which 'cloud computing' relies require vast amounts of energy and water to remain at an operational temperature[12]. In the short term, reducing this cost has obvious

[10] Adamatzky, A. and De Lacy Costello, B. (2002), 'Experimental logic gates in a reaction-diffusion medium: The XOR gate and beyond', *Physics Review* 66(4) [http://link.aps.org/doi/10.1103/PhysRevE.66.046112].
[11] This process is essentially that described by Adelman in 1994, although other researchers have since generated innovations in technique (Adleman, L. (1994), 'Molecular computation of solutions to combinatorial problems', *Science* 266 pp.1021–1024).
[12] An article for cleantech.com suggests that, on average, between 40 and 60 per cent of the energy costs of a data centre are claimed by cooling mechanisms (Rick Cockrell, 'Putting

commercial benefits; in the long term, these data centres will not be sustainable, as energy costs rise and water grows more precious. Biological computing may also be more sustainable as the components and raw materials necessary are in good supply and relatively easy to obtain, unlike many minerals used in the manufacture of conventional computers.

Computers built along non-silicon lines offer the promise of greater computing power and with it the potential of managing some of the difficult logistical problems facing us (air traffic control, global just-in-time supply chains, ensuring secure electronic communications, to give three examples). But the real future impact of this development in computer science is likely to be more far-reaching. Technology has always given us metaphors for describing the wider world and the action of our minds[13]. For the ancients, their hydraulic technologies led the Greeks to describe the soul as *pneuma* and Galen to imagine four humours; in the late 18th century it was clockwork that underpinned the workings of the world and the rational movements of the mind; it was steam power that built the repressed society diagnosed by Freud, and the new telegraph that helped Helmholtz convey his ideas on neural connectivity. In the present day, the ubiquity of the computer places it within easy grasp when we reach for a way of describing the flow of goods around the world or the ways decisions are taken, and make it easy for us to imagine that a process or action can be imagined as a flowchart, a series of steps to be taken in sequence.

But the new forms of computer that we see emerging today break with this tradition in important ways, and if they become the new mechanical metaphor we might expect to see some new cognitive habits take hold across society. Where we currently imagine processes to be linear and sequential, we might soon conceive of them as parallel and stochastic. Alongside the understanding that a single test-tube might hold trillions of computers we might see our sense of scale change, moving away from the human scale on which our rooms and furniture are constructed to a greater fluency with the very big and the very small. Aware that the laws of probability and statistics play a part in the success of these computers, we might move away from an assumption that particular outcomes inescapably follow particular starting conditions, and towards a new acceptance of the uncertainties that underpin life. “If X, then Y, on the whole” might become our new default attitude.

---

the heat on data centre cooling costs’, 28/09/2009 [http://cleantech.com/news/5063/putting-heat-data-center-cooling-co].

[13] Rodney Brooks goes into more detail on this tendency in his 2008 Edge article, ‘Computation as the Ultimate Metaphor’ [http://www.edge.org/q2008/q08_5.html#brooks]. John Daugman gives an overview of the various technological metaphors employed by the cognitive sciences in ‘Brain Metaphor and Brain Theory’, in Schwartz (ed.), *Computational Neuroscience*, MIT Press, 1993

Implications for education

The immediate implication for education is that learning institutions will have access to levels of computing power far beyond those available today though possibly not in the form of individual desktop computers. Current trends towards ‘cloud’ and ‘utility’ computing, in which data storage and access are provided to a user by a third party, might mean that educational access to computing power is ensured through university partnerships (with local institutions’ degree of access reflecting the relationship they have with the university) or through commercial computing providers. In either case, the massively parallel nature of biological computing might mean that the issues encountered in the 1960s and 1970s when sharing a limited resource (time on a central computer) can be avoided. This access to computational resources would enable institutions to carry out many more activities: students might replicate physics experiments carried out in the late 20th century (a time when the necessary computing power was only available to a few researchers). Students could design and run simulations and models of natural phenomena which are vastly more complex than anything to be found currently in Second Life or similar virtual worlds. School productions (film or stage) might be augmented with digital effects: schools themselves might be augmented with real-time information generated by their students’ activities. The ‘disappearing classroom’ or ‘virtual school’ that has been such a feature of debate for the past few years might finally make a (non-) appearance. In the most general terms, educators will have to be aware that those calculations and symbolic manipulations which they currently imagine to demand an inaccessible level of computing power will become trivial for their students once they enter the workforce.

In order to support the growth of this new form of computing, demand will grow for educators to ensure that bioscience, genomics and psychology are a substantial part of any future curriculum, particularly in connection with engineering and computer science. Biological concepts and understanding of the complex behaviour of biologically-derived networks will be needed not just by those building these new computers but by anyone who hoped to understand the effect they have in a future society.

Conclusion

To many readers, these brief discussions of technology trends might feel familiar. But the value of considering them again lies in the opportunity to reconsider where we think their impact will lie: for the most part, discussions of technological innovation tend to focus on the external effects we expect to see – this operation will be N times faster, or that opportunity will be available to X million more people. This is only part of the story.

What all of these trends have in common is that they threaten the boundaries we currently use to make sense of the world. In everyday life, we have certain understandings around the phenomena we are familiar with – “this starts here, not there”, “that is the task of these groups but not of those groups”, “this is relevant in these contexts but not in those contexts”, and so on. We have these understandings not only for education, for health, for communications technology, display technology, family life, state involvement in individual lives, work, our relationship with government, but also for deeper concepts such as the difference between mechanical and biological life, for example, what is alive and what is not alive, what constitutes a causal relationship, what the relationship between agency and action is, how far our perceptual experiences are shared, how far it is possible to imagine another human’s inner state of mind. These are fundamental building blocks of our lived experience, unquestioned in daily life and central to our ability to operate in the world. But all of these are challenged by the developments charted briefly above: trying to hang on to the boundaries we use currently to navigate the world will make it harder for us to adapt when it changes.

Most crucially, then, these changes challenge how we currently understand ourselves to be human. You don’t have to be a believer in the Singularity or transhumanism in order to understand that in the present day we operate with an understanding of the boundaries between our inner lives and the external world, and that technological advances alter how relevant these definitions are. Perhaps the most challenging change of all will be learning to understand that someone in the same space as us may be experiencing a very different reality: just because we are both looking at an object in the same space, there’s no reason any longer to assume that we see the same thing. This will move from being a philosophical and academic question to being a practical issue in our everyday lives. We’ll cope, of course: look at the way we’re evolving expectations around the use of mobile voice technology, with different groups seeing shared norms and rules of behaviour emerge around when it’s polite to let people know you’re taking a call, when it’s better not to, when it’s good to share music with the rest of the bus and when it isn’t. We won’t necessarily do everything that technology lets us do. But we can’t help but have our behaviour shaped by it.

So there are some psychological and social changes ahead for us to navigate (as there have been many times before), although the nature of a complex world makes it impossible to say exactly what these might be. Education, then, needs not only to respond to the specific demands of these trends – more prominence for concepts from the biological sciences, a greater preparedness for interdisciplinary working – but also to ready learners for uncertainty, equipping them with the qualities they’ll need in order to respond best to the changes they’ll see in their lives. To succeed in a society facing complex and uncertain changes, people will have to be psychologically resilient, conceptually agile and emotionally robust. Whatever the future has in store for us, teachers who can best foster these qualities will be as vital then as they are today.