

Review

Technology and the future of healthcare

Harold Thimbleby

College of Science, Swansea University, UK

Significance for public health

Technology drives healthcare more than any other force, and in the future it will continue to develop in dramatic ways. While we can glimpse and debate the details of future trends in healthcare, we need to be clear about the drivers so we can align with them and actively work to ensure the best outcomes for society as a whole.

Abstract

Healthcare changes dramatically because of technological developments, from anesthetics and antibiotics to magnetic resonance imaging scanners and radiotherapy. Future technological innovation is going to keep transforming healthcare, yet while technologies (new drugs and treatments, new devices, new social media support for healthcare, etc) will drive innovation, human factors will remain one of the stable limitations of breakthroughs. No predictions can satisfy everybody; instead, this article explores fragments of the future to see how to think more clearly about how to get where we want to go.

Introduction

Pluck a nurse and surgeon out of the nineteenth century and transport them into a modern 21st century hospital and it would be a thoroughly recognizable place, with the same hierarchies and strict cultures. Patients treated as helpless, stripped of their clothes and possessions, lying in beds and almost completely ignorant of their illness. They might be disappointed in our treatment particularly of old people, but I don't think it would surprise them.

If our two time-travellers were able to attend a post-mortem and listen in on a discussion of human error, very little would seem novel. Clinicians would still be in denial, lawyers would still be hovering, and the *delay and deny* culture would be no surprise. However, the changes that would surprise the nurse and surgeon are all changes to technology. Infusion pumps, dialysis machines, antibiotics, heart valves, MRI scanners, even hand washing stations would be new ideas. All the hidden technology used in the laboratories behind the scenes, from path labs to decontamination, would be startlingly new if it was noticed.

Although the medical culture is similar, there have been dramatic technological changes, and actually these changes would be hard to explain. Does anybody even know how an infusion pump works? They used to be clockwork (and before that, gravity fed) and now almost everything contains a computer and has a colourful screen and lots of buttons. Implanted defibrillators that use telephone networks and web sites to keep cardiologists up to date with their patients are just magic; new pharmaceuticals that change moods, change blood pressure, or kill bacteria: all are modern magic. On reflection, given the centuries

of stability, it is amazing how much healthcare has changed in the last 150 years – and one wonders how this accelerating pace of change will proceed in the future.

Arthur C. Clarke, the prolific futurist and science fiction writer, famously said that any sufficiently advanced technology is indistinguishable from magic. Perhaps the main difference for the couple lifted out of the nineteenth century is they are sure it's magic, whereas we have stopped thinking about it, and just take it for granted!

Time-traveling fiction starts to dig into many interesting issues we might normally avoid thinking about. What about organ harvesting? What about enhanced humans? What about the end of antibiotics? Will robots take over? And so on.

Some of what seems to us today like science fiction is going to be routine in the future, perhaps even in our lifetimes. Yet much of today's human story about relationships, hopes, error, grief and denial is going to remain entirely recognizable in the future. We will still have authority gradients, we will still have controversy over human error, and patients will still be made helpless so that they are easier to treat. The reason is that technology is driven by the market: if somebody has an idea that they can turn into a physical realization that they can sell, they can also patent it or license it, and thereby make a return on their investment. This, in turn, will encourage them to find ways of making it smaller and cheaper, and marketing it on a larger scale; thus it is technology-driven. In contrast, human culture does not make profit for anybody. Improving culture means admitting somebody or some process was not good enough to start with, and who wants to do that, especially when lawyers are watching? There is little economic incentive to improve culture.

Atul Gawande makes a similar point about the World Health Organization's *Surgical Safety Checklist*.¹ The checklist is a simple piece of paper that helps change human behaviour – it helps check the patient has been correctly identified, helps introduce everybody by name (to reduce authority gradients), and so on. It's simple and saves lives – more effectively than many drugs. Gawande asks who will promote such an idea when nobody makes any profit from it? It's just a bit of paper anybody can print. If it was a technology (such as a patented drug) that promised the same improvements in outcomes, everybody would be buying it, and the pharmaceutical company making it would be promoting it heavily. Patients would ask for it to be used. But a piece of paper anybody can print is not exciting enough. Crucially, the only person who benefits from the checklist is the patient (the clinicians benefit indirectly, because more successful operations mean less litigation). The patient is probably unconscious at the very moment they ought to be asking for it!

Science fiction

Our time-travelling fiction is a small example of the power of using science fiction to help envisage and plan our future. In contrast to the usual tunnel vision prediction of future trends, which often highlight glowingly positive ideas, science fiction lets us explore *and communicate* futures we want to live in by telling rounded stories we can engage with. More importantly, science fiction can also explore

dystopian futures we want to avoid; George Orwell's 1984 helped avoid his dystopia happening (so far anyway).

In the present paper we have not space to create further stories, but we commend the method to both manufacturers and consumers of technology – the hospitals, clinicians and patient groups, and especially to designers.²

When we tell good stories, we get into them, but there is not one story about the future. Anything is possible, and we need lots of stories, to explore good and bad and indifferent choices. Moreover, when we get to the future, it too will have another future. There is not one future, but many. We will never find satisfactory solutions to anything, as there will always be new things to try and explore. This week it might be the Cloud or improved natural language processing, but before we've got that working properly, somebody will have invented something that solves even more problems and sounds equally seductive. Yet while technology drives changes in healthcare, the fundamental problems of wellbeing, health and happiness, will remain.

The easy story is that the future will be better. Technology will advance and there will always be new and exciting solutions. Today we have robotic keyhole surgery, and things can only get better. We have intelligent decision aids to improve diagnosis, and they will only get better. Some people would point to the underlying drivers: technology is getting faster, better and smaller. Moore's Law says the speed of innovation is accelerating. The simple story is we will just enjoy the ride. However, the more complex story exposes trade-offs. For example, new computers are indeed much faster, but to take advantage of them we first have had to throw away the slower computers so they can be replaced, and then we may well discover the patient information on the old computers won't work on the new ones. In fact, we are in a continual struggle to keep up – it costs us a lot, and a lot of *solutions* that excited us yesterday are already in landfill. The faster we go, then, the more we can expect incompatibilities, and indeed greater spread between those at the sharp end of developments and those without the resources to benefit. When we honestly think about the future, we have to broaden our spotlight from the few exciting ideas that attract our attention to the wider issues, the broader context of change and complexity, in which those innovations could be used effectively. As good science fiction does so well, turning an exciting idea into a fully-worked out story helps us explore the issues more realistically.

Rather than develop a single story about the future, this paper now turns to presenting principles, themes and scenarios that a good writer might integrate to create a coherent picture.

Key points about futures for healthcare

Patients are the reason for healthcare and they should be at the centre of it. This article, however, is about possible technological trends and drivers in healthcare; it should therefore be read in conjunction with patient-centered perspectives like the Royal College of Physician's Future Hospital: Caring For Medical Patients report.³

Technology does not have an agenda of helping healthcare, however much we might like to focus on benefits. It develops because of miniaturization, lowering costs of production, and so on, not because it makes people well, but rather because it can find ways of making money and reinvesting it. Koppel and Gordon's edited book *First Do Less Harm* is recommended as an overview of issues.⁴

The pace of change is accelerating: our time travellers from a hundred years ago were surprised at a few things, but had they travelled back in time, with a few minor bumps like William Harvey discovering blood circulation, very little changes all the way back to Hippocrates.

Human nature does not change, at least not on these technological timescales. The authority structures in healthcare, the division of labour, the pretence that clinicians know everything, and other human factors are slow to change. Despite our knowledge of germ theory and antisepsis, we are still resistant to washing our hands.

There are many futures to plan for. As soon as we get to our future, there will be another – and we will increasingly be seeing partially-completed solutions superseded by even better ideas. Today we might be thinking we *just* need to computerize all patient records, but before we've finished doing that some fancy new technology will change what we want to do or how we should do it. For the foreseeable future, we will have to live with fragmented and partially working technologies.

We need to take the future seriously as, literally, it is all we have, and certainly all our children will have – and we can be certain that as we get older, we are going to end up with all the problems of old age. Surely, we want healthcare to improve in the future? We should put effort into future planning, not once, but continuously.

In the rest of this article, we will make repeated comparisons between technological factors and human factors. Often these drivers are not aligned (one of the key messages of this article), and technology is therefore unlikely to develop in ways that are optimal for healthcare on its own accord. Our concepts of ourselves, from conception to death, as individuals, families and as communities, are inextricably linked to technological possibilities. We ought to have a *Future Healthcare Institute*, which will be kept continually busy prioritizing and reprioritizing principles to guide and align healthcare and technological developments together. One imagines such an institute giving guidance legal and regulatory guidance, for example as has already happened in *ad hoc* ways in some countries addressing advances such as fertilization technologies.

Technical factors

Healthcare is just a market for technology where consumers such as hospitals are happy to pay enormous amounts of money, particularly for prestige equipment, such as PET and MRI scanners and linear accelerators.

Accelerated cost savings

Technology automates and extends things that previously had to be done by people. Before infusion pumps, nurses had to give injections every so often; the infusion pump technology automated that. Now the nurse's time is freed up for other activities, and if the manufacturer has used technology in the production of the infusion pump – as they surely will have – they can reduce the cost of production for exactly the same reasons. Some plastic moulding process will make millions of infusion pumps as easily as it makes one; once one infusion pump has been programmed in software, it costs essentially nothing to program them all. This virtuous circle of using technology to make technology ensures prices drop, market share increases, and profit margins increase, which in turn allows the manufacturer to invest in more cunning production and distribution technologies.

However, what is important to notice is that these benefits do not accrue to custom or rare problems that cannot be mass-produced. This means that a technology like an MRI scanner that can scan anyone equally well is going to be much more popular than a technology that has to be customized to a particular patient's conditions.

Personal healthcare

Already, the assumptions of mass production are changing. For example, today's *3D printers* are capable of making objects of any shape; they are slightly less efficient than standard mass production, but the costs of custom objects of certain sorts has been lowered significantly. It is now possible to custom make titanium implants the right shape and size to fit. Going further, it is widely envisaged that custom drugs will be manufactured, customized to the patient's disease and

genetic makeup. While this seems to be enormously beneficial to patients, there are dangers. For example, a customized drug may be very effective, but its side effects will be unique to the patient too, and therefore harder to diagnose and manage.

Personal healthcare has an interesting technological imperative. If we can personalize healthcare, we get population-sized markets: instead of selling to clinicians, manufacturers can sell to individuals – a market 1000s of times larger.

Big data

Patients generate huge amounts of information – *patient records* – from X-rays to blood test results. Replacing paper with computerized summaries makes patient care easier and more efficient. In the future the quantity of information will increase dramatically because of genomics (and the huge genomics of our symbiotic bacteria) and personalized medicine, and as more patient data is collected, more insights will become available.

If computers collect data on patient illness, treatments and outcomes, one automatically obtains valuable information on the effectiveness of those treatments, or relations between side effects and patient characteristics across whole populations. Huge amounts of data will be collected, hence the name big data. Once the infrastructures have been set up, the incremental cost of adding one new patient will be essentially nothing, and this economy of scale will drive further technical developments. Epidemiologists will benefit enormously, but the benefits to individuals are less obvious, except in the long run from big data's contribution to the progress of medical science more generally.

Social media, patient power, mobile health and education

Stopping people going to hospital in the first place and empowering people to care for themselves and their families is something computers are already doing well. But as patients are empowered, is their newfound knowledge helpful or unrealistically raising their expectations? Today, the internet is problematic, as there is no consistent way anyone can distinguish snake oil from sense, how patients can distinguish reasonable treatment from misguided hope – there will always be a lot of solutions hunting for the patient's money. Technical solutions to this problem include providing accredited high-quality information; cultural solutions include improving education. When somebody has a knee injury at 40 this should not be the first time they encounter the bewildering amount of variable information and social media on the internet! Their management of their condition – whatever it is – would be much improved if they had been exposed to sensible strategies since preschool.

Dramatic, transformational integration of technologies

There is not space here to fully explore the vast range of likely and significant technological breakthroughs. Consider nanohealth, brain implants, artificial organs, networked sensors, genomics, exoskeletons ... just a few of the potentially transformative developments already under way. Some of these technologies are going to transform our whole approach to illness and health – in the same way that the nineteenth century development of anaesthetics changed society's moral approach to pain. Pain and suffering used to be inevitable; now we like to think we have a right to painless procedures – and in turn this has influenced everything, from our treatment of patients to our treatment of animals (*why should animals suffer?* is a very modern question). New technologies, like nanohealth, are going to have ethical implications that will be hard to anticipate. Sometimes ethical issues will be hard to negotiate because they will be apparent only after somebody has got things working and already has a business-driven perspective.

Security, privacy and monitoring

In a world beset with major security concerns (like terrorism) it is inevitable that all technologies, even in those healthcare, will be aligned with national priorities. For example, taking patients' fingerprints and other biomedical identifiers will become easier (perhaps driven by consumer finance, such as credit card security); and, as it becomes easier, gathering data for state security will happen as a side-effect of routine clinical practice. The state will be able to identify illegal immigrants and outlaws and others; the current notion of *patient confidentiality* will be eroded in a way that will be impossible for clinicians to control.

Today we may think this would be objectionable, but it is salutary to remember that we happily divulge all sorts of personal information during our use of mobile phones, credit cards, as well as during our use of the internet. We unthinkingly sacrifice our privacy because of the huge convenience of buying stuff on the internet. It seems to make losing our identities a trivial price to pay. When considering future healthcare trends we can expect similar trade-offs; it will be easy to slide into levels of surveillance we do not now like, falling for it because of the healthcare benefits we want. Surveillance is not the only downside of course – paying data rights owners; paying software licenses; signing off responsibilities for insurance liabilities – all happen, and are often signed off without sufficient thought.

It is increasingly trivial to collect data about patients and the quality of patient care. This information can be aggregated and help discover variation in treatment and outcomes, and hence help improve quality – which is good. On the other hand, data inevitably distances the manager from the patient as an individual: perhaps the fundamental notions of patient care will lose out to organizational or state concerns, because cost management and security, not care, becomes to be the point of the information.

Health 2.0

There are many areas where the scale and unit profits of the healthcare market will drive technical developments. Collectively, this technology-driven progress in healthcare is sometimes called *Health 2.0*, to distinguish it from what we are doing now – Health 1.0. While Health 2.0 is exciting, it is sobering to realize that perhaps it is just the start of an upgrade path: Health 2.0 will have problems we solve with Health 3.1, and in turn that will develop into Health 4 or whatever. While it seems obvious technology will continually advance, it is going to be harder to ensure that each iteration of technology satisfactorily achieves what it claims to achieve, without having to be fixed up and upgraded soon after.

Unfortunately, few manufacturers stay in business selling us perfect solutions; they stay in business by selling us something to keep us consuming: a service, something to rent, a disposable product, a product that wears out, or a product that goes obsolete. Certainly Health 2.0 will lead inexorably to more developments, whatever they will be. The danger is that it will make us eager to upgrade before we have even realized the promised benefits of Health 2.0. Somehow, we need to work with manufacturers to align their interests of staying in business with our interests of having a predictable and stable life. We might do that by distinguishing *infrastructure*, which is provided about once, with *consumables* that are provided regularly. This is the economic model of infusion pumps: you buy an infusion pump once, but the giving sets are replaced after each infusion. Over time, the manufacturer makes more profit on the easily reproduced plastic tubing than the complex pump, and everyone is happy.

In some areas, the *consumables* will be information itself. This costs nothing to reproduce, but people own it and want to make a return on their investment. Thus patient data will be *owned* so that its owners – rarely the patients! – can make money from it. Information is stored in

computers in data formats, and often these are proprietary: the format of a patient data system belongs to the manufacturer. This leads to the danger that the patient data is inaccessible except on the terms the manufacturer imposes. It may be costly to convert it into other formats, say to upgrade to a different manufacturer's systems. Worse, if a manufacturer goes bust, some data may be lost. This is a very real problem, as our inability to use data on paper tape, cards, cassette tapes, magnetic tape, VHS tapes – none of them very old technologies – and so forth, testifies. A desirable technological trend, then, in fact a trend that bucks the trend to date, has to be the assurance that data remains accessible and usable over long periods of time – at least a 100 years, which is way longer than any electronic technology!

Hacking and open health

Healthcare sensors can be readily bought off the internet, and it is easy for technically-minded people today to build sophisticated equipment (to *hack*) to collect and analyse any personal or clinical data using their own computers. Credit-card sized computers like Arduinos and a few biomedical sensors cost about the same as a drug prescription!

Some individuals are already obsessed with collecting as much health-related data as they possibly can about themselves – it is not just people will illnesses, but people who want to lead healthier lifestyles or be better athletes. If these people upload their data and contribute to aggregated data, they are contributing to citizen health – just like *open science*,⁵ except tackling healthcare problems. At its simplest, they would be contributing to epidemiological studies; at its best, they would be helping build databases and web systems that other people can find their medical conditions in, and hence find support communities. Many patients end up with more time on their hands than they expected, and this is how some choose to use their time: solving their own problems and helping others.

Hacking is not restricted to patients: a doctor using a laryngoscope has the choice of paying commercial prices for a video recorder (*e.g.*, to record images to send to an ENT specialist), or more conveniently recording the video on their iPhone – about a 1,000 times cheaper.

The point is, technology is empowering people to do what they want to do, and in the future patients are going to take some of the initiative away from professional healthcare, particularly for diagnosis, chronic illnesses, and lifestyle advice.

Technology is diverse and surprising

These are some of some powerful technological drivers, and it is hard to draw a line under the discussion. We have not discussed many technologies that are both critical and exciting such as nanohealth, personalized healthcare, mobile health, telehealth and so on – the beginnings of all of these are already available and in use in first adopter places. What the brief discussion illustrates is the diversity, the rapid pervasiveness, and the complex trade-offs of future technologies.

Bridging future gaps

The science fiction author William Gibson says the future has already happened – we just don't know where. All the ideas we discuss in this article about the future have happened.

From considering technological drivers, we now turn to *human* futures. We believe these will be more stable and less likely to change, but will raise increasingly unexpected interactions with the new technologies. In areas like human error this is alarming, for if we believe that technology improves – why else would we adopt it? – then, as human error will still occur, the error must even more surely be due to the humans involved, not the *improved* technologies which were intended to design out error. In other words, the irresistible drive to adopt improved technology may exacerbate our management of human error. The economic drivers that push technologies have vested inter-

ests in promoting benefits and belittling problems. And healthcare has no end of problems: we all want and expect better care, costs are rising and performance is declining; living longer, and living with chronic illness, are other problems. Healthcare staff are over-worked and under-resourced... it is hard to imagine technology changing that. On the contrary, many technologies (take MRI scanners, heart implants) are very expensive, and buying into them will exacerbate financial pressures.

Safety, security and regulation

In the future there will remain an enduring distinction between *safety* and *security*. In healthcare these mean different things: safety is about patient and staff safety – basically, following Hippocrates *first do no harm* – and security is about controlling access, in particular so that intruders, rogue patients and staff cannot get inappropriate patient access, whether that is informational access or physical access.

Security means stopping bad people doing bad things. If a bank loses money to fraud, this is not unexpected – we all know there are plenty of bad people around who want to get at our money. It follows that it is the bank's responsibility to provide security.

Safety means stopping good people doing bad things. If a nurse is involved in an untoward incident, this is neither normal nor expected. It is easy, then, to think the *good* nurse has gone *bad* and therefore they are to blame – this is the conventional *bad apple* approach to safety. Indeed, if a good nurse has gone bad, this is a serious betrayal of our high regard of the nurse, which makes things even worse. The bad apple theory is very appealing: getting rid of this *bad* nurse appears to solve the problem.

In short: security is seen as an organizational responsibility (*e.g.*, the bank's or the hospital's), whereas safety is seen as the individual's responsibility (*e.g.*, the nurse's). Technology improves things that generate return on investment (security, speed, efficiency, scale and reach) and safety will not do that while users are scapegoated.

Moreover, safety is hard to assess up-front, unlike simple claims for low price, speed or efficiency. Unless regulation requires safety to be assured, we would expect safety to take second place. We therefore anticipate an increasing debate between safety concerns on the one hand and *regulatory burden* on the other. Since currently the regulatory burden for technology is negligible, certainly compared to the rigors of pharmaceutical development, much could be gained by strengthening regulation. We suggest careful attention needs to be paid to statutory regulation. To avoid hasty regulation that is ineffective or rapidly obsolete, we need to think very clearly. Today there is a lively debate about regulating computer technology; some say (for example) mobile apps should be more tightly regulated; others say that rigorous protocols (such as randomized controlled trials) take so long the technologies will be obsolete once there is formal evidence one way or the other.

Solving the right problems

Conventional patient records are paper records in folders in cabinets. Many patients have extensive patient records, lab results and so on, and even more patients have patient records that are in many places – in hospitals they have visited, consultant's offices, general practices, and so on. They are rarely all together where the patient is, often they get lost or duplicated, and sometimes destroyed by fire or floods. Many healthcare providers have trucks shipping patient records around their areas.

The obvious thing to do is to computerize all the records, and then use networks to ensure they are always available wherever they are needed. Looking at records on a screen is simpler than wading through piles of paper. Since computers already work, all we need to do is set up a program to scan or type up all the existing paper records. Job done!

Unfortunately this obvious solution creates new problems.

When a clinician examines a patient, they want to refer to the rele-

vant parts of the patient's medical history. If we have simply computerized the patient records, all we have done is made the large, scattered piles of paper into something that can be viewed on a computer screen, but now the clinician can only view one window at a time, and they may easily lose the big picture. Information may be scrolled off the screen, or be concealed behind pop-ups. In fact, we have merely swapped the unusability of piles of paper for the unusability of a user interface.

While we are very familiar with the ways that paper records can fail, unfortunately we are much less familiar with the ways that computerized records are hard to use and may mislead us.

Tom Landauer's book *The Trouble with Computers* makes the insightful point that computers are enormously successful in areas where there is commonality that can be computerized — banking, communications, stock control, office documents, email are all good examples.⁶ Areas where they have been less successful are those where success depends on the human element. My bank account works just like your bank account, so computerizing either of our accounts is the same as computerizing everyone's. But my patient records are different to yours. The computerization of my records does not help computerize yours or anyone else's. Well, that is not quite true. Computerizing my records helps computerize yours, but when those records are used, we and the healthcare professionals using them will have different problems. As the healthcare computer systems scale up to handle more patients, the usability problems get compounded — in contrast, as bank accounts are scaled up, things become more uniform and easier to automate successfully. (Banks also have a very different approach to problems; a British bank does not have to handle my Russian currency or it can charge me exorbitant rates, but a hospital that ignored my X rays would be negligent.)

In healthcare, we have to pay more attention to the broad context of how information is used.

This is the concern of User Centered Design (UCD), which Landauer describes well. The ideas have been taken up in international standards. In particular, healthcare technologies must be developed using UCD processes, such as ISO Standard 62366 etc.⁷ One of the most important features of the standards is that they make clear that new technologies are not going to be perfect and need to be tested and improved to better match how people actually use them.

UCD is essential in the battle against information overload and the *law of unintended consequences*. Technology is introduced to solve a problem or to improve performance, but this then changes people's behaviour and new problems may emerge.

Originally, email seemed like a wonderful idea — it is cheap, fast, saves paper, and so on. But we are victims of its very success: now people have so many emails that they are overloaded (it is hard to prioritize), to say nothing of spam and phishing, flames and people sending irrelevant or erroneous emails to thousands of recipients. It is now possible for an ill-conceived email to waste thousands of hours when it is sent to many staff. Emails are a recognized and growing problem; but the same trend is affecting test results, patient records, drug-drug interaction reports. For all of these reasonable tasks it seems obvious they should be computerized, but doing so often results in increasing amounts of low-level information that can distract people from doing their *real job*.

UCD helps because it emphasizes that no innovation is ever finished: we have to see how it is used, and continually improve it. Email, and the rest, have a way to go, and UCD promotes that at each step we should be user-centred (driven by the needs of users and what they are trying to do) rather than technology-centred.

Unfortunately, technology creates new users. Computers need technicians and managers, and these users also contribute to the UCD improvement cycle. However if we are not very careful, the management of the technology gets a life of its own that takes a higher priority than delivering improved patient care. When investments are made,

the experts are consulted — but now the experts appear to be the technologists rather than the healthcare professionals or even the patients. This can cause many problems.

Systems that are under-performing and hence need improving often induce workarounds by their users. For example, passwords may not work very well, so nurses find ways to get on with their jobs regardless. Unfortunately the people the other side of the computers just see the systems apparently working; they do not see the workarounds or the unintended risks nurses may be creating as they get things to work. When the system is improved, the workarounds are not considered sufficiently, and the new system may have unanticipated problems that even workarounds cannot overcome.

The science (rigor) dream (unqualified success) gap

X-rays were discovered by Wilhelm Röntgen in 1895 and immediately recognized as having huge potential for healthcare. Only a few years later, one of Thomas Edison's assistants, Clarence Dally, who had been enamoured with the potential of X-rays died, of cancer because he had been experimenting with them every day (Figure 1).

It is now obvious that X-rays are not risk-free. Every exposure to X-rays helps a patient yet at the same time exposes them to risk; it is now routine to make a careful trade-off between the benefits and risks. Similarly, we now recognize that pharmaceuticals are not magic and risk-free. In fact, we hardly understand how many pharmaceuticals work, and it is routine — in fact, a requirement — to perform the *gold standard* randomized control trial (RCT) and other forms of careful experiment before allowing drugs to be released to the market for wider use. Despite our best endeavours, we have a growing awareness of worrying and complex side-effects, such as growing antibiotic resistance that has arisen from over-enthusiastic use of antibiotics (not least in animal husbandry). Some of the original miracle antibiotics are no longer effective.

Ben Goldacre's provocative book *Bad Pharma* documents how pharmaceutical development sometimes side steps best practice, and promotes drugs that have untoward side effects and other problems.⁸ For example in a large RCT, some patients may die for unrelated reasons.




Figure 1. Clarence Dally X-raying his hand, from the *New York World*, August 3, 1903, page 1 (American Newspaper Repository).

The question then is what to do with their data; it is very tempting to treat them as if they had survived and had been cured. Another example success bias in the scientific literature: authors of scientific papers want to publish their successes rather than their failures. So the literature under-represents drug trials that fail or uncover unwanted complications. In turn, this means that systematic studies of drug trials cannot get the correct baseline for experiments, since many experiments are not published. Regulation is starting to address this problem.

Goldacre makes clear that pharmaceutical development fails scientific standards; yet technology development, such as robotics or computer system development, does not even aspire to the scientific standards that pharmaceutical research is aware it fails to reach. As Goldacre puts it, pharmaceutical R&D is close to immoral; yet technology R&D is ignorant. This is worrying for the future, as technological developments may not be adequately tested, tested without satisfactory controls, and under conditions of vested interests. (Most of the studies arguing Physician Order Entry systems are either good or bad are written by people using the single systems they are publishing about; they are not good science.) New technologies may have unfortunate side-effects or other problems (such as forcing risky workarounds) that nobody has seriously looked for, let alone rigorously assessed.

The last paragraph makes a strong claim, but it is justifiable. Modern infusion pumps will have been certified for clinical use, and thus evidently passing the applicable safety tests and standards. Yet modern infusion pumps are driven by computer software (*e.g.*, in their firmware), which can be modified at will by the manufacturers and parameterized by hospital technicians. Indeed, software is regularly updated to fix bugs and to make minor tweaks. Modifying software can completely change the behaviour of devices. What makes the software control of devices so appealing is that manufacturers can create a variety of devices for different market sectors all on top of the same architecture. Changing the software can change a device from, say, a simple infusion pump to a dose error reduction *intelligent* pump. But such changes can be made after it has been certified for use, without any further regulatory control. Furthermore, one will struggle in vain to find any scientific literature on the assessment, let alone RCTs, for such devices. It does not exist. On the contrary there is a growing literature on the safety problems of infusion pumps.⁹

One can hope that the future trends include tightening the culture of technology development. The standards need improving, and the *laissez faire* culture of contemporary development needs addressing. In many ways, today's medical technology is at a maturity level comparable to the *unsafe at any speed* automobile of the 1960s.¹⁰ If we do not address the lax scientific standards of technology development, it is unlikely that future technology interventions will improve safety or other desirable measures of healthcare performance.

Human factors

The nature of human expertise is that it makes errors likely,¹¹ and clinicians are highly-skilled experts. To become expert at some process means automating it, doing some or all of the task without continual reference to the wider situation. For example, when you learn to drive a car, you are consciously aware of many factors (such as clutch control), but as you gain expertise, driving becomes automated and you are able focus on higher-level goals. As an expert driver, you may find it seems easy to hold a conversation on a mobile phone, as you now have the spare cognitive resources to do so. Unfortunately if something unusual happens, say if a child runs into the road, you may not be paying enough attention to the situation to take appropriate action – ironically, when you were less of an expert driver, you would have had to pay very close attention to road conditions, and you may not have been driv-

ing so fast either! The point is that as new technologies will improve things, we humans will still make errors.

Human factors is already a problem today: complicated gadgets induce use errors. There is a balance between the time and effort a manufacturer is going to spend making some technology *easy to use* (and *safe to use*) when the economics of selling the product may not prioritize those qualities. Usually new technology is procured because it promises to improve efficiency or reduce costs; safety is an issue covered by insurance, and is rarely part of the procurement requirements. Regulation requires some basic safety, but the features that sell devices often conspire to make the use of the devices more complex. As we discussed with modes, above, the more features the more complex a device is to use. Yet typically features sell technology, and the difficulty of use is pushed onto problems for the users. If there is an untoward incident, it is much easier to blame inadequate training (*i.e.*, the user's incompetence) rather than the complexity of the system being used.

Human factors – issues such as situational awareness, tunnel vision, and so on – is a large and important area. There are two questions for the future: how can technology help, and how can technology be improved to be intrinsically safer?

How can technology help?

In time-pressured environments, humans often suffer from tunnel vision – focusing on the original task and overlooking larger situational awareness. The classic example is intubating a patient. This is a demanding, time-critical procedure. The longer it takes, the more pressurized the clinician is to complete the procedure. Sometimes the patient will get into problems, and a tracheotomy is needed – urgently. Sometimes the clinician is so focused on the intubation that the warning signs are missed, with disastrous consequences. Here, technology can help by using monitoring technologies. Remote monitoring of the patient's vital signs can enable a trained person out of the pressured environment to make helpful comments: *Hi! I can see you have about ten seconds before you need to stop...* Such ideas lead to concepts like the *electronic ICU* (eICU) where a control room monitors perhaps 100 beds. Experience with such technologies has been positive, especially if careful steps are taken to avoid a *them and us* division between the clinicians doing the work and the clinicians monitoring them. Usually a staff rotation is used, so everyone experiences both sides of the camera.

How can technology be designed better?

Manufacturers can use better design processes, such as those outlined in standards such as ISO 62366.⁷ Doing this effectively is hard work, and with *rush to market* it is tempting to do the bare minimum, for all the reasons discussed elsewhere in this article. Here we can mention three useful ideas:

Discovering use errors takes a long time, and this conflicts with rapid entry to market. The solution may be to design systems so that they can be improved in the field. This is actually easy – firmware is routinely upgraded for bug fixes anyway. What needs doing is logging device use in sufficient detail so that the manufacturer gets a good insight into how the device is being used or is failing to be used. Currently, this information rarely gets back to manufacturers in a useful form.

Many use errors follow predictable patterns. So-called *post-completion errors* are common and hard to eliminate just by improving human procedures. A nurse may use a blood glucometer to measure a patient's blood sugar levels. The nurse moves on to the next patient, and then puts the blood glucometer in a docking station to upload all readings. As described, at this point the nurse has failed to make a written note of the levels on the patient's paper notes, yet because the device has docked, it has discarded all recordings. This is a post-completion error: the nurse made the error after they had finished. The solution is to

redesign the technology, and there are many options here. Why delete uploaded data, for instance? Why not have a reminder in the device to confirm the nurse has recorded data before taking another reading? Why record manually on paper patient notes anyway? This is an example of how the standard operating procedures combined with built-in technological assumptions induce errors (which in this case are unprofessional, and perhaps disciplinary offences), but more thoughtful design can avoid them.

Physicality

The huge leverage computer technology brings, because it is *virtual* and can do anything with information (and hence the same piece of technology can be mass produced for a huge market that has not been pre-conceived), has a down side. Humans are physical.

The problem can be illustrated very simply. In the *old days* books where physical objects, and they looked and felt different. A well-read book would look worn, and an unread book would look new. You would recognize bookmarks sticking out of books, you could write annotations in the margins, you would know how much you have got left to read before you finish it. You could put a book by the front door of your house to remind you to pick it up in the morning; you could leave a book by your bed so it was ready for next time you wanted to read yourself to sleep. And so on. Now, with electronic readers, all books and documents look the same – like the general-purpose computer they are on. Of course, the computer can create colourful covers and images, but the physical object is always the same: namely, the computer or tablet.

In the *old days* a patient would go to their doctor and get a paper prescription. They would then go round to the pharmacist and get their medicines. One problem with this process was that paper prescriptions were notoriously hard to read, and there was a danger of incorrect dispensing. Today, this process has been computerized. The doctor sends the prescription electronically to the pharmacy, and the pharmacy can dispense the drugs almost immediately.

Unfortunately, the patient has lost the physical prescription in this process. They leave the doctor's appointment carrying nothing. There is *nothing* to remind them to go to the pharmacy to pick up the drugs. Indeed, pharmacies are now having to dispose of drugs that were *efficiently* dispensed as soon as they were prescribed but were never collected by the patient.

Just as electronic books are a gift to book publishers – because the expensive paper book is replaced with a cheap electronic book, free to reproduce once just one copy has been prepared – increasing use of computers in healthcare is irresistible. X-ray photographs no longer need to be developed, put in folders, held up against viewing screens. They can be emailed. But what they have gained, they have lost in physicality. It may no longer be possible to put an X-ray on a viewer back-to-front (hence causing a left/right error), but it is very easy to look at the wrong patient's X-ray – because they all look the same without their physicality.

Unfortunately physicality is in direct conflict with the driving economics. Pieces of paper are very familiar and we understand exactly how they work, but computer screens are cost-effective precisely because they can display lots of information repeatedly without addition costs. Yet from a safety perspective, the screens all look so similar they may be confused.

One solution to this is *skeuomorphism*: making new technologies imitate old (and hence more familiar) technologies, otherwise known as using the right design metaphor. The example here would be to improve the display so that the *paper* information being displayed looks more like actual paper – perhaps with torn edges, discolouring if frequently used, and so on.

A second solution is *affordance*: make things look like how they are intended to be used. For example, a cup with a handle has the affordance that encourages you to pick it up by its handle. Particularly in

emergencies, people need to know what to do *intuitively* – and affordance is a key part of design.

Thirdly, we can make computers *disappear* – what is left is a physical object which happens to do something complicated, but it has been made invisible to the user. RFID tags and smart chips can be embedded inside objects so that they can do things and interact with each other, but the object appears to be *normal*. A nice example is the hotel key card; so far as a hotel guest is concerned they behave just like room keys, but inside some sophisticated cryptography creates a host of benefits – for example, unlike losing a physical key it is not a problem if you lose a key card, and the locks do not need to be replaced. When patients routinely have ID tags embedded in them, many of the current problems of patient barcodes will disappear – no doubt to be replaced by different problems!

This article is not the place for a deep discussion on design,¹² but it is important to point out that looking up skeuomorphism or affordance on the internet reveals a huge variety of conflicting opinions. In other words, these ideas are ways to *argue about* design, not ways to design. One still has to do detailed work to make future technologies successful, regardless of their supposed affordances. Despite having nice words to talk about good design, no future will be an inevitable success.

Enhanced and bionic humans

Although we have said that human factors will remain a constant for the foreseeable future, there is nothing to stop technology enhancing people. Technology will not be used just to monitor and help make people well, but will be used to make them better – better as sportsmen, as healthcare professionals, or whatever. We already have augmented reality, where surgeons can see an enhanced view of the insides of patients. If patients can have brain implants to enhance their lives (to manage Parkinson, for instance) surgeons will have implants to improve their skills, using robotics to reduce tremor or computers to reduce error. After all, the core of the humble calculator – which reduces drug dose calculation errors – has already shrunk in size from table-top clockwork of 50 years ago to something so small it could be swallowed.

Who are the protagonists?

Thinking about the future is science fiction. Key to the success of any story is to help the reader identify with the key character; or in the present context of thinking about healthcare futures, are we writing from a patient's perspective or from a technologist's, or a healthcare professional's, or ... the story makes choices, and if the choices align with our own interests the story seems more plausible and persuasive.

If the market has developed profitable technologies, it follows that we (you and me) want those technologies. This truism needs emphasizing. Industry stays in business making what we want to buy. Industry is adept at adapting to make what we find irresistible: this is market competition. The manufacturers who are better at seducing us survive and grow. As a consumer, I love iPads (at least I did in 2013), but that does not mean that iPads can do much good in a professional healthcare environment. Should we fill hospitals with iPads? One part of me, the private consumer, says yes, they are wonderful! It is interesting how Blackberries were driven by employers, but iPhones and iPads are being driven by consumers; increasingly the employee is dictating what technology the organization they work for uses. But the other part of me, the clinician, the scientist, the technician, asks for evidence that they will truly improve healthcare. It doesn't make such a good story for me to say *I am not sure!* If I am in healthcare procurement, I have to think very carefully to distinguish what I covet as a consumer from what will actually improve the organization I work for.

It is certain that getting evidence, doing experiments (RCTs?) will delay the market, and there will be huge pressures on us to believe in the wonderful visions, rather than in the necessity of careful development. Healthcare is complex, and just throwing technology at it will not in itself change anything usefully, other than costing us a lot of money (which is exactly what the market wants us to do). As we approach the future, we need to learn how to plan our resources much more carefully and experimentally. The subtle cost-benefits of X-rays were not apparent immediately, and they rather ruined the original magical story.

People trying to sell us the future will certainly latch into our acquisitive consumer natures – this is who we are as individuals. They are less likely to latch into us as patients or clinicians, and thus – while being an engrossing story – they may miss the point and sell us on technological ideas that do not really improve the world so much as satisfy our consumerist urges. One of the important things science fiction teaches us is that the future will not be populated only by sensible solutions. While we hope for a happy ending, there will be problems and even wicked plans, evil empires and natural disasters to overcome along the way.

Who makes the future?

The future of healthcare is about the patient (or stopping people becoming patients), but patients are not the main stakeholders in healthcare. Insurance companies, big pharma, doctors, managers, suppliers, builders, governments and many other forces will influence the future. Will innovation help patients or will it be partly to help monitor clinicians delivering healthcare?

It is interesting that since hospitals started introducing computers that the ratio of managers to clinicians has steadily increased; it is not totally clear that computers have made delivering healthcare more efficient or safer, but they have certainly increased the volume and turnover of the business. Moreover, now patient records are computerized, with the obvious benefits, there are also problems. To get problems fixed, upgrades have to be purchased, and this can be very expensive. Patient records were once on paper; now they are in proprietary formats, and the manufacturer can *lock in* the user to their particular system, so buying upgrades – perpetuating the lock in – seems cheaper than moving over to an alternative system.

Conclusions: what we need to do

If we don't know where we are going, we won't know when we don't get there says the quotable Yogi Berra. The market will surely figure out a way to make money, and technology will advance in miraculous ways. Instead, we need to figure out a way to have healthier and happier patients (and not just treating them as individuals), and to do that we've got to focus on integrating technology with culture rather than focusing on technology alone.

In my view the key thing is to think clearly. Initiatives like the UK Royal College of Physicians Future Hospital Commission, where thought-leaders – patients and healthcare professionals – have sat down and worked out what sort of future hospital they need are to be commended;³ more importantly, they have articulated principles, not just specific solutions. Such principles are or should be timeless; we should not plan the future by being technology-driven (*e.g.*, implementing *cloud*, *nanohhealth*, etc) but by improving along criteria behind principles (such as improving patient care or staff support). Articulating

the principles of the futures we want should be a continual process, not a one-off activity; every day there is a new future to plan, and new discoveries that will change our minds about what is possible and likely. Future planning should be as much a routine part of healthcare as responsive care is. If we don't know what we need, we will get what is easy and profitable to make; as we emphasized above, what we need and what we want are often confused.

Are we making the healthcare future we want? If people from the past suddenly materialized in front of us, we would want to be able to convince them the world is the better place they had hoped for. What stories would they take back to help direct their technological trends and developments?

Correspondence: Harold Thimbleby, College of Science, Swansea University, SA2 8PP Swansea, Wales, UK. Tel. +44.179.220.5678.

E-mail: harold@thimbleby.net

Key words: future healthcare technology, human factors.

Acknowledgements: the author acknowledges the very stimulating and encouraging comments most especially from Annegret Hannawa, Ross Koppel, Don Norman, Ben Shneiderman and the anonymous referees.

Conflict of interests: the author declares no potential conflict of interests.

Funding: this work was funded by the UK EPSRC on grant EP/G059063/1.

Received for publication: 1 November 2013.

Accepted for publication: 1 November 2013.

©Copyright H. Thimbleby, 2013

Licensee PAGEPress, Italy

Journal of Public Health Research 2013; 2:e28

doi:10.4081/jphr.2013.e28

This work is licensed under a Creative Commons Attribution NonCommercial 3.0 License (CC BY-NC 3.0).

References

- Gawande A. The checklist manifesto: how to get things right. New York: Picador USA; 2011.
- Johnson BD. Secret science fiction. *Computer* 2013;46:105-7.
- Future Hospital Commission. Future hospital: caring for medical patients. September 2013. Available from: http://www.rcplondon.ac.uk/sites/default/files/future-hospital-commission-report_0.pdf.
- Koppel R, Gordon S, eds. First, do less harm: confronting the inconvenient problems of patient safety (the culture and politics of health care work). Ithaca: ILR Press; 2012.
- Nielsen M. Reinventing discovery: the new era of networked. Princeton: Princeton University Press; 2011.
- Landauer T. The trouble with computers: usefulness, usability, and productivity. Cambridge, London: The MIT Press; 1995.
- Usability Engineering ISO 62366:2008 - Medical Devices 2008.
- Goldacre B. Bad pharma: how drug companies mislead doctors and harm patients. 4th Ed. London: Faber & Faber; 2012.
- Thimbleby H. Improving safety in medical devices and systems. Proceedings IEEE International Conference on Healthcare Informatics (ICHI 2013). September 2013. Available from: <http://harold.thimbleby.net>.
- Nader R. Unsafe at any speed. New York: Pocket Books; 1966.
- Reason J. Human error. Cambridge: Cambridge University Press; 1990.
- Norman DN. The design of everyday things. New York: Basic Books; 1995.