GRESHAM COLLEGE

4th June 2020.

Digital Healthcare: Will the Robot See You Now?

Professor Richard Harvey FBCS

Before preparing this lecture, I spent some time listening to the excellent Gresham series on Medical Education by Roger Kneebone. In particular, his lecture on "The Ethics of Surgical Innovation" [1]. His opening quote was the rather striking phrase "Technological developments in surgery take place a dizzying speed," yet my lecture was going to open with observation that, to me, technological progress in surgery seems slow and technological progress in the rest of medicine is even slower! Obviously, Professor Kneebone is an expert and he can see progress that the rest of us do not register. Part of the reason is shown in Figure 1 which displays data from the World Health Organization (WHO).

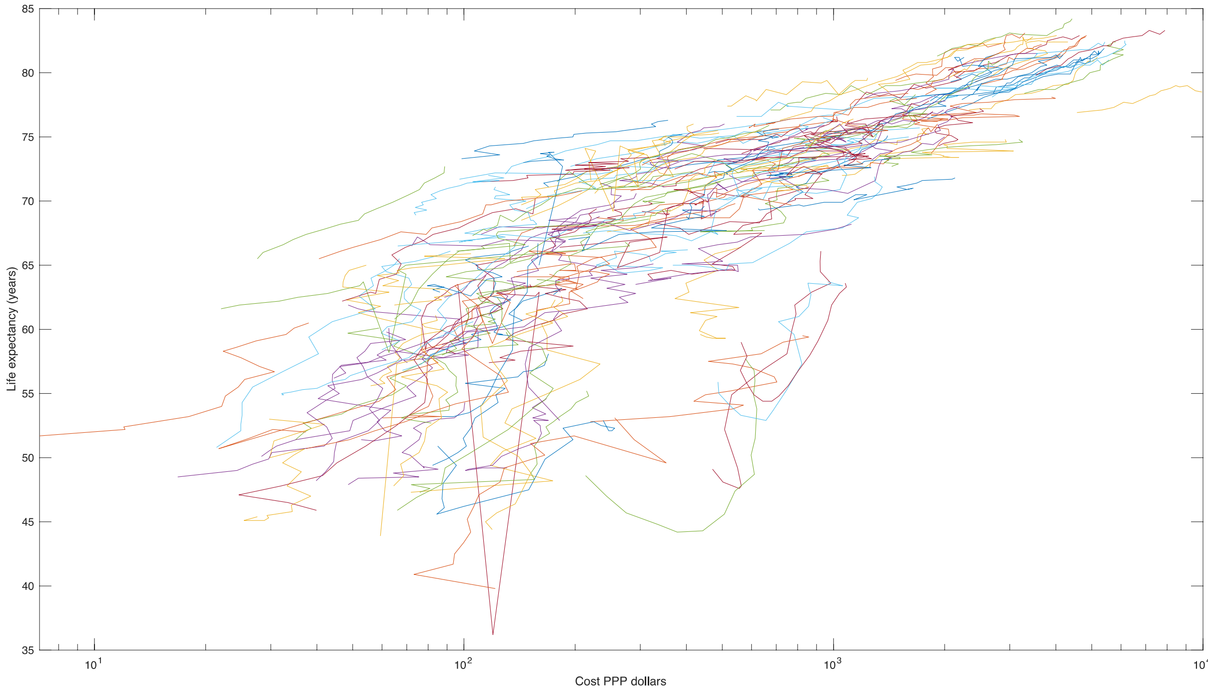

Figure 1: *Life expectancy (both sexes) from year 2000 to 2016, for 193 countries, versus the cost of healthcare per person in purchasing power US dollars.*

Most of the lines show a shallow increase in life expectancy[1] which of course is very welcome, but it is only really visible over a long timescale. Another evident feature is that, apart from countries which have strife, pretty much all countries are increasing life expectancy. Countries in the Global South showing more rapid increases than those

[1] I would estimate the gradient to be between 0.5 year per year and 1 year per year in developed countries with no warfare.

in the Global North. The outlier on the right with lower than expected life-expectancy and staggeringly high costs (note the *x*-axis has a logarithmic scale) is the USA. The UK is tucked in the leading countries and is neither spectacularly cheap nor spectacularly effective. In the "premier league" for life expectancy are Japan, Switzerland, Spain, Singapore, Australia, France, Canada and Italy.

However, it's tricky writing about healthcare as it seems to be a highly politicized issue. For some, the very idea of making profit out of healthcare is an anathema (even when profits get reinvested). For others, the State is so creaky and unreliable that it simply cannot be trusted with people's lives. And strangely, two countries which have a lot of shared cultural values, the USA and the UK, violently disagree on healthcare provision. Even more strangely, many countries which are thought of as quite socialist countries have extensive involvement of the private sector. This lecture is not about the politics of healthcare but, as we shall see, even quite modest technological innovation can get caught-up in polarized arguments. This in itself is a disgrace, as there are many technological innovations which are being stifled by luddite practices, undue caution and unreasonable behavior. If there is a lesson in this lecture for a British audience it is that unquestioning acceptance of historic practices of the National Health Service, which is Britain's health system, is dangerous and is killing people. If there is a lesson for the US audience, then it is the outrageous costs of US healthcare ought to be contained, not exacerbated, by investment in new technology. For other audiences, be happy that you are not in a country with health-system wars!

That said, this lecture was originally scheduled to be given, in person, in London. But the British government issued a "stay-at-home" order, the lecture was postponed, so I've given this lecture a re-write to reflect the new role that IT has had in the COVID crisis.

I've structured this lecture in terms of the usual classification of health interactions: *quaternary* which are highly experimental surgeries, research trials and so on, *tertiary* which is routine but complex activities typically confined to big hospitals, *secondary* which are the specialists delivering routine care usually in a hospital and *primary* which is general practice or Primary Care Providers (PCPs). I've introduced a new category which I'm designating *self-care* which is patients looking after themselves. Given that most of the development activity in IT is funded by the private sector there is an implicit calculation which is the product of market size with sales price with a factor than can be difficult to define but is something to do with regulatory effort. To put some numbers on it, we can consider NHS England[2]. England has population of around 63M people. They make around 300M visits to GPs per year. Hospitals in England manage around 6M visits per annum. So, an app that works for the general population has a market size ten times larger than one for hospital visits. Hence the two sectors that look immediately appealing to IT companies are self-care (huge market, little regulation but low sales price) and quaternary care (huge sales price, little regulation but low volume). The intermediate markets are more challenging because they are relatively small and have regulatory environments. These environments also lead to a paucity of suppliers with associated worries about profiteering and monopolies.

Quaternary care is very popular in the media since developments in the lab are the future and the future is a popular subject with the press. Furthermore, experimental healthcare is usually only possible in the most seriously ill patients so there is much human drama which is the meat and drink of journalists. Computing developments in quaternary care is a whole series of lectures in itself but two current themes are Artificial Intelligence (AI) and robotics. Artificial Intelligence has been covered in several of my previous lectures.[3] One application that recently made the media was the use of deep neural networks to identify new forms of antibiotics. In robotics we should mention the recent arrival of Google Health.

However, it is self-care and primary care where we can see the most activity. Apple and Android both have health interfaces and there is much commercial interest in connecting your phone or smart-watch to cardiac monitors, blood pressure measurements and so on. Apps are particularly well suited to types of condition where it is very important to keep on track with routine measurements, Type 1 diabetes for example, or where there is an

2 For those unfamiliar with the United Kingdom I should point out that each country of the UK has its own variant of the NHS. I cannot find any logical technical reason for this, but it certainly allows more politicians to have a go at being a Health Minister.

3 And those of Yorick Wilkes who has been talking about AI and its societal consequences.

implication of extensive talking therapy (depression for example) or where there is an environmental cause, asthma for example. The NHS health apps library lists hundreds of such apps and one can imagine a future in which your physician prescribes an App for your condition[4].

But what about secondary and tertiary care? Here we are entering a very difficult area. I can see at least three factors which are guaranteed to stifle innovation.

First is the observation that IT systems for hospitals have a track-record of being delivered late of not at all. The most notorious was the NHS patient record system [2]. After around £12 Bn spend, the project was abandoned[5]. Clearly it is brave manager that agrees to take on a health IT transformation project. Furthermore, if they do take it on, it is so much more remunerative to be working for the IT contractor than for the health service so there is inevitable leakage of effective staff to the private sector.

Second is the legitimate concern that software that is used in clinical settings can have bugs [3]. The most notorious example is known as Therac-25 and is now a standard case history in courses on software engineering. Therac-25 was a computer-controlled radio-therapy machine produced by Atomic Energy of Canada Ltd. Older versions of the machine had mechanical interlocks to prevent dangerous radiation doses being delivered to patients. The Therac-25 removed these and replaced them with software interlocks written in PDP-11 assembly language. Figure 2 shows an example of one the simplest programs one could write in MACRO-11, which is the PDP-11 programming language. I challenge you to understand the program let alone to spot any errors!

```
.TITLE  HELLO WORLD
.MCALL  .TTYOUT,.EXIT
HELLO:: MOV     #MSG,R1  ;STARTING ADDRESS OF STRING
1$:     MOVB    (R1)+,R0 ;FETCH NEXT CHARACTER
BEQ     DONE     ;IF ZERO, EXIT LOOP
.TTYOUT          ;OTHERWISE PRINT IT
BR      1$       ;REPEAT LOOP
DONE:   .EXIT

MSG:    .ASCIZ /Hello, world!/
.END    HELLO
```

Figure 2: “Hello world!” program in MACRO-11 taken from Wikipedia (https://en.wikipedia.org/wiki/MACRO-11)

The programming for the Therac-25 was done by one programmer whose identity was never revealed. There is no evidence that he or she was qualified and little of the code appears to have been documented. The machine was notorious for throwing errors and one of the operators reported around 40 errors a day. To anyone experienced in software design, these factors alone would be alarming and unsurprisingly the interlocks did not work properly. Between June 1985 and January 1987, at least six patients were exposed to massive overdoses of radiation [4] with associated injury and death.

Astonishingly, given that virtually every software engineering course in the world discussed the Therac-25 case, it happened again in 2001. A treatment planning system sold by Multidata Systems, miscalculated the dose for

[4] There is also an NHS App which was designed to give UK patients access to booking appointments and to repeat prescriptions. It was launched in December 2018 and described by the UK Health Minister as a “world first”. It wasn’t, and the app is widely reviled for its poor usability and extraordinarily finickity enrolment procedure.
[5] I reckon that to be around £500 per UK taxpayer, a quite staggering sum.

radiation treatments and five people were killed[6]. In [3] are presented a litany of poor practices and dangerous assumptions which imply that errors in medical computer systems are highly prevalent.

Third is regulation. An obvious reaction to the first two factors is to press for standards, inspections and regulations. These are expensive to meet, so health IT projects become even slower to implement, meaning technology moves on faster than the project using it, and of course cost rises.

Yet Time-and-Motion studies of clinicians, [5] for example, give highly alarming figures for the amount of time spent wrestling with EHR (Electronic Health Record) systems. In [5] it was reported that 52% of a working day is spent on EHR tasks of which the majority was documentation, order entry, billing and coding. In [6] it is noted that nurses spend around 37% of their time with patients despite it being known that the time spent with patients is associated with improved patient outcomes. Of course, healthcare is not alone in the scandalous amount of time required of professionals to fill in clerical information but surely the purpose of IT should be to minimize that time?

It may seem that we have created a logical impasse or dilemma in which the parts of the health system which need the most improvement are the least likely to receive it. In practice, I believe this is resolvable. Firstly, there are fourteen recommendations in [3] many of which are proactive and do not amount to simply more regulation. It would be good to take action on those. Secondly, while these doubts about medical software remain, it should be relegated to roles that are less likely to harm patients. Thirdly there are a number of standard taxonomies of systems which match the development effort to the amount of testing required. In defence and commercial projects the Technology Readiness Level (TRL) scale is a common one: the idea is that the complexity of each system is assessed from TRL Level 1, which is basic research, through to TRL 9 which is the system operationally deployed. Although originally developed by NASA to help characterize technological readiness for space-flight, the system has adapted well to different environments. A corollary is that the higher the TRL, the more comprehensive should be the testing and evaluation. This idea has been re-purposed for a medical environment with the Stead scale [7] which is illustrated in Table 1. The idea illustrated in Table 1 is that horizontal arrows indicate an improvement in evaluation complexity and vertical arrows indicate technology improvement. Thus, a system is firstly specified, and that specification is evaluated in the definition phase. As the technology is developed it moves from bench trials and field trials (undertaken by the inventors) and then further trials (undertaken by independent bodies).

[6] The radiotherapy technicians were sent to jail although it seems evident that the machine had calculation errors.

Relationship of System Development Stage to Level of Evaluation*

System Development	Evaluation I Definition	II Laboratory Bench	III Laboratory Field	IV Remote Field Validity	V Remote Field Efficacy
A Specification	→	↓			
B Component development		↓			
C Combination of components into a system		↓→	↓→	↓	
D Integration of system into environment		→	↓→	↓→	↓
E Routine use			→	→	→

*The rows indicate stages of system development and the columns represent the different levels of evaluation. The presence of an arrow in a column of a row indicates that the level of evaluation indicated by the column is appropriate for the stage of development represented by the row. A horizontal arrow indicates that it is appropriate to proceed to the next level of evaluation while the system development stage is unchanged. A vertical arrow indicates that it is appropriate to proceed to the next stage of development without changing evaluation level. A double arrow indicates that it is appropriate to proceed to either the next stage of development or the next level of evaluation.

Table 1: *The Stead scale, taken from [7].*

There is a whole lecture to be written about methods for software development[7] and to avoid this lecture being a litany of faults with medical IT systems, I'd like to turn to some more positive developments.

The first is telemedicine, also known as telehealth and eHealth. Telemedicine has a long history and the commonly cited first example is that of the transmission of radiographic images over a distance of 24 miles by telephone in 1948. I think one can go earlier than that, and the Wellcome trust archives contain an advertisement dating from 1850 for a gutta percha speaking tube so that sleepy GPs can converse with night visits from the sick and their messengers without the danger and inconvenience of encountering the cold night air[8]. Nowadays there are two modes of telemedicine that are of interest. The first mode is inter-health-system. This is now so routine that it is hardly worth mentioning – clinicians routinely consult other clinicians by all the usual means. However, as a frequent visitor to hospitals I should record my frustration at waiting for a highly paid expert to stride from one part of the mega-complex to another so that he or she can ask me exactly the same questions which have already been asked by five lesser qualified brethren. Surely it cannot be beyond the wit of hospital IT directors to have a video calling system by the bedside so that a ten-minute consultation does not expand to a thirty-minute absence[9].

The more interesting opportunity is extra-health-system in which clinicians are in contact with patients who are at home or work. In the UK, the best known of these providers is Babylon Health although they also operate in Rwanda and China. There are several aspects to the Babylon Health app, it offers video consultations with physical doctors, it has some symptom checkers (Clinical Decision Support Systems in the jargon) and it offers an electronic pharmacy service. The reception to Babylon is mixed: Health Ministers and Mandarins have been enthusiastic;

[7] The next series of Gresham IT lectures will definitely have at least one lecture on software development.
[8] The advertisement contains some less than convincing testimonials from surgeons and doctors.
[9] This lecture was prepared during the covid-19 outbreak and I note that the use of telepresence robots, which has been routine in education and business for quite a few years, has been remarkably stimulated in Asian hospitals.

Doctors less so possibly because Babylon have a penchant for inflammatory remarks including the claim that their system is better at GP Exams than real GPs [8].

What is undoubtedly true is that there is quite a bit of Venture Fund money in symptom checkers, chatbots and AI. In the lecture I identify at least $1Bn of investment in start-ups and were we to extend that to large companies such as IBM Watson (estimated investment $1Bn) and Google Health then the sum would be very large indeed. The large companies have tended to steer away from computer diagnosis and look for niches. In IBM Watson's case it was cancer. There is a very large literature on cancer, it is changing all the time and oncologists can find the rate of change overwhelming. So, the idea was to train Watson to read the oncology literature and recommend treatments. The concept is that clinicians can cross-check their treatments with their learned computer friend in the corner. Again, the outcome is highly contested, health care managers seem positive; doctors are skeptical [9]. The other large IT company to be involved is Google. Google had several prominent initiatives in healthcare including a new approach to robotic surgery in conjunction with Johnson and Johnson. Robotic Surgery is in tertiary care so I would judge is well-suited to Google's experimental and innovative style. They have also partnered with the Royal Free Hospital on an App called "Streams." Streams aims to alert carers to imminent and dangerous conditions in hospital patients. The initial application considered Acute Kidney Injury management although it is thought the it might be suitable for other conditions, such sepsis, where speed of response is critical. As with all the other AI applications, the introduction of Streams was controversial: that notorious busybody the Information Commissioner ruled that Royal Free London had breached data protection regulations by releasing 1.6M patient records to Google Deep Mind without appropriate legal frameworks and safeguards. Such rulings may be correct in law, but they create a rather chilling atmosphere for machine learning in decision support, at least in the UK, since machine learning works best, and is safest, when it is able to train and test on very many data points.

Notwithstanding all the criticisms of decision-support AI, and the politicization of health care, and the intervention of the Information Commissioner, my own view is that AI as a decision support tool is highly likely in future medicine. For the foreseeable future it is highly unlikely to replace doctors, but it is highly likely to deployed in ways that either improve safety, lower cost or reduce admin for the medical profession.

While AI is rather glamorous and the subject of much dinner-party speculation, there are also computerization activities which are much more firmly focused on running-cost reduction; improvement of patient experience and returning doctors and nurses to the front line. The grandest of these are digital hospital projects. These are popular in all the leading countries in healthcare and in the lecture, I have videos illustrating Humber River Hospital in Canada and New Karolinska Hospital in Sweden but I could have picked St Stephens in Australia, the Royal Adelaide Hospital in Australia and probably quite a few in Singapore. Brand new buildings are very impressive, and the builders are rightly proud of the digital innovation not only in the hospital but also in its construction. However, health care is best thought of a system – the hospital might be a hub but social services, primary health care, social care and pharmacies all need to function together to be seamless and effective.

This leads to the thorny problem of Electronic Health Records. The UK is often held up as a model for health research because of the wide availability of health records. A patient in the UK might therefore reasonably assume that the information available to their GP is also available to the emergency medic in the ER, (or A&E as it is called in the UK). They might also assume that the data from that expensive private consultation might also be fed into their health record. Unfortunately, this assumption is so far from the truth that it is almost laughable. A recent paper [10] computed the probability that a patient that had treatment in adjacent government areas would have their records held in the same type of EHR. For most regions in the UK, that probability is zero, so Britain has sick patients trundling around the UK with bundles of paper.

Health systems therefore present a very peculiar landscape to the IT Professional: there are huge numbers of interesting research projects using AI, Big Data, robots and so on and yet the basic information architecture of record-keeping, security, privacy, safe-design and so on, with the exception of a few leading countries, is primitive and fragmented. For taxpayers this is a familiar but expensive picture: money is easily spent on high-visibility projects such as robot doctors but the large bills, which relate to fixing the information architecture, are yet to come. However, there are a couple of areas which do not represent huge spend and may payback very tidily. The first is education. Medical education is very expensive and suffers from the problem that mistakes can kill people. Mistakes can be reduced by "problem-based learning" in which students spend time on the job rather than in theoretical lectures but that is even more expensive. The solution, copying from the aviation industry, are simulators. At the moment the focus is on impressive animatronic simulators but "phantoms[10]" and software-only simulators are also key. Simulation and the introduction of new roles such as Physician Associates might relieve some of the pressure caused by increasing demand. But what if demand itself could be reduced?

Leaving aside the rather obvious remark that healthier living ought to lead to fewer healthcare visits (and there is a barrage of Apps devoted to various theories on healthy living), there is evidence that many hospital interventions are simply not required by patients [11]. It appears that when patients are properly briefed on the consequences of some medical interventions, they decide to not take part. In some cases, the number of operations halves. Of course, one might argue that doctors and nurses should jolly well be briefing patients accurately but it's a complex business – the "right" decision depends on a patient's personal preferences, outlook and so on. Discovering these takes time and in the UK a GP consultation lasts ten-minutes. In other settings, ten-minutes is just about long enough to discuss the weather, the football and make a cup of tea. So good future practice will be for your primary healthcare professional to prescribe you some decision-support software so you can consider all the options in your own time.

All the above was written before the COVID-19 crisis and remains valid. But what is the role of IT in the current situation?

Firstly, it is worth applauding the IT industry for superb planning and design – huge numbers of people have to dramatically change the way they worked and lived from the physical to the virtual. And what has happened to the internet and telecommunications infrastructure during that change? Nothing[11]. Unlike the supermarkets who despite being given advanced notice, ran out of toilet paper; the internet has not run out of bits, IT suppliers have not run out of cameras and 5G has continued to supply high quality signals.

When it comes to the pandemic response, the contributions are patchier and much more experimental but broadly fall into three categories: pandemic prediction, disease prediction, contact tracing. Notable in pandemic prediction is BlueDot, a Canadian company formed to spot the early signs of pandemics. Automated systems scour the grey literature looking for signs of unknown diseases, they then couple that with predictions about air travel[12] and hence produce early predictions of the spread of disease. BlueDot's claim is that they described the COVID-19 problems earlier than the WHO which, if true is an impressive achievement. Disease prediction is a fascinating area but much more nascent. The idea is that by continuous or near continuous monitoring your body's signals (heartbeat, oxygen saturation and so on) that technology might know you have Covid-19 before you yourself know. This would be a major advance and it is not fanciful – there are claims that the monitoring of wearables can predict Lyme's disease which is a notoriously tricky disease to diagnose.

However, the vast amount of media attention is concerned with contract tracing. The idea behind contract tracing is that, if we know who is infected then we isolate those people and the disease cannot propagate or rather the mass propagation of the disease at the same time is prevented[13]. The most dramatic of contact tracing systems in democracies appears to be South Korea. Here the government tracked infected citizens mobiles phones using mast triangulation[14], they tracked purchase cards for the same reason and hey presto!, with what would appear to be a gross invasion of privacy, they have an effective contact tracing system. Such a system does not require citizens to have a smart mobile phone. The alternatives do require citizens to have a smart mobile phone which

10 A phantom is an object that is designed to look like a human under, say, ultrasound, x-ray or other imaging modality.
11 Well Netflix had to increase the compression on some of their movies, but I suspect that was largely about constraining cost.
12 Air travel is a major contributor to the spread of pandemics.
13 It is something of a bugbear of mine that critics of lockdowns constantly quote the relatively low death rate of the disease. But that death rate is on the assumption that a health care system can cope. If we all get the disease at the same time, then the health system is overwhelmed, and we will be back to medieval practices for a few months. It is that situation which western governments want to avoid. In countries with very underdeveloped health systems where only the richest can afford treatment, lockdowns might not be the ethical response.
14 See my Gresham Lecture on Crime for an example of this in action.

means that models suggest that is required to be an 80% adoption of any contact tracing app. That is a daunting number[15]. The basic idea is the same for all the apps. When a citizen becomes infected, they press a button and everyone they have been in contact with gets notified and, hopefully, goes in to quarantine. Where they differ is in whole holds the list of citizens. The British NHS and Poland have settled on a system with a centralized database which brings with it the inevitable concerns about security, practicality and so on. Whereas other countries prefer a decentralized approach in within, either, the central authority is only notified when the button is pressed or there is no central authority, your phone stores an anonymized record of who it has been in contact with, when you press the button through a natty bit of cryptography you contact everyone without you or them knowing who you are. Leaving aside technicalities, the real issue is how to get 80% adoption. E H Fowler when writing about split infinitives, wrote that were a group of people who neither know not care what a split-infinitive is. Likewise, there are, I suspect, a huge number of people who neither care not understand the implications of a central versus the decentralized system. They will be happy with any system that works. We can count on them to adopt the system. Then there are people like me who know, but no care (very much). We will adopt. But what about the people who care very much? Even if they do not understand the issues. They will not adopt the app. And we need 80% to adopt the app. That is likely to be a problem.

I hope this has been a useful romp through the digital healthcare landscape. Compared to other lectures in the series, this one has had a very unusual flavor. It is commonplace for healthcare professionals to describe healthcare as complex, multifaceted, nuanced and so on. Some of that complexity has appeared in this lecture for which I apologize, however the root cause of digital healthcare is quite simple. Healthcare is very expensive, so only governments or very large organizations can provide system-level coverage, and, with the exception of a few leading nations, the information architecture is far too fragmented to provide confident care. That fragmentation is caused by a failure to design and invest in basic information elements such as databases, networks, security, training and so on. Certainly, in the UK if the NHS was compared to a house then it would be a house with leaky plumbing, dangerous wiring and a roof full of holes. But there would be marvelous people living in that house, full of care and concern. There would be a Ferrari in the garage and a 4K TV in the living room, but everyday functions would be a struggle because of the poor infrastructure.

From my perspective that is a pretty undesirable situation and we should hope for dramatic improvements. If you live in country where government controls healthcare, then I would advise to press your government to think really hard about healthcare as a system. IT Professionals can design improvements to systems, they can plan and anticipate demand but a hodge-podge of political add-ons, initiatives administered by very complicated management structures and ad hoc political interference, that is an accident waiting to happen.

1. Roger Kneebone and Chris Peters, The Ethics of Surgical Innovation, Gresham College Lecture, https://www.gresham.ac.uk/lectures-and-events/trust-surgical-innovation
2. “NHS IT system one of worst fiascos ever”, BBC News 18th sept 2013, https://www.bbc.co.uk/news/uk-politics-24130684
3. “Computer bugs in hospitals: a new killer”, Martyn Thomas and Harold Thimbleby, Gresham College Lecture, Feb 2018, https://www.gresham.ac.uk/lectures-and-events/computer-bugs-in-hospitals-a-new-killer
4. N. G. Leveson and C. S. Turner, "An investigation of the Therac-25 accidents," in *Computer*, vol. 26, no. 7, pp. 18-41, July 1993.
5. Arndt, B. G., Beasley, J. W., Watkinson, M. D., Temte, J. L., Tuan, W. J., Sinsky, C. A., & Gilchrist, V. J. (2017). Tethered to the EHR: Primary Care Physician Workload Assessment Using EHR Event Log Data and Time-Motion Observations. *Annals of family medicine*, *15*(5), 419–426. https://doi.org/10.1370/afm.2121
6. Westbrook, J.I., Duffield, C., Li, L. *et al.* How much time do nurses have for patients? a longitudinal study quantifying hospital nurses' patterns of task time distribution and interactions with health professionals. *BMC Health Serv Res* **11,** 319 (2011). https://doi.org/10.1186/1472-6963-11-319
7. William W. Stead, MD, R. Brian Haynes, MD, Sherrilynne Fuller, PhD, Charles P. Friedman, PhD, Larry E. Travis, PhD, J. Robert Beck, MD, Carol H. Fenichel, PhD, B. Chandrasekaran, PhD, Bruce G.

[15] Lower adoption levels are more likely which means one cannot rely on the app alone.

Buchanan, PhD, Enrique E. Abola, PhD, MaryEllen C. Sievert, PhD, Reed M. Gardner, PhD, Judith Messerle, Conrade C. Jaffe, MD, William R. Pearson, PhD, Robert M. Abarbanel, MD, PhD, Designing Medical Informatics Research and Library—Resource Projects to Increase What Is Learned, *Journal of the American Medical Informatics Association*, Volume 1, Issue 1, January 1994, Pages 28–33, https://doi.org/10.1136/jamia.1994.95236134

8. "Babylon claims its chatbot beats GPs at medical exam," Jen Copestake, BBC Click, 27th June 2018, https://www.bbc.co.uk/news/technology-44635134
9. How IBM Watson Overpromised and Underdelivered on AI Healthcare, Eliza Strickland, IEEE Spectrum, 02 April 2019, https://spectrum.ieee.org/biomedical/diagnostics/how-ibm-watson-overpromised-and-underdelivered-on-ai-health-care
10. Warren LR, Clarke J, Arora S, et al. Improving data sharing between acute hospitals in England: an overview of health record system distribution and retrospective observational analysis of inter-hospital transitions of care. BMJ Open 2019;**9**:e031637. https://bmjopen.bmj.com/content/9/12/e031637
11. Patients' preferences matter; stop the silent misdiagnosis, Al Mulley, Chris Trimble, Glyn Elwyn, The King's Fund 2012, ISBN 978 1 85717 637 7