Patient Safety

Stories for a digital world

HAROLD THIMBLEBY

with

PRUE THIMBLEBY

CONTENTS

QUESTIONS

INFORMATION

TRUE STORIES

FICTIONALISED TRUE STORIES

THINK LIST

CAT THINKING

One day, we were eating a meal at the kitchen table. We could hear Po, our new kitten, scrabbling around in his litter tray, scattering litter to try to cover his poo. He proudly walked into the kitchen and jumped up on our table. Po was an infection risk! His feet had been paddling in poo, and now he was strutting around the kitchen table. "Get off!" we shouted at him.

Po rolled onto his back, and purred.

As if on command, we tickled his tummy. Aaaaah. He's so sweet...

In less than a second, we went from being sensible people, worrying about infection risks and food hygiene, to mindless people seduced by a furry, purring kitten.

We just forgot all about hygiene. He was so cute!

...

Driving our uncritical happiness was a cocktail of hormones: endorphins, dopamine, oxytocin, norepinephrine, and prolactin, as we stroked and petted the little kitten.

Our brains, driven by hormones, left us no choice but to feel good. Po purred at the centre of our attention. The world was a happy place! And we both totally forgot all about the bugs at the end of his sharp claws.

Here's the insight: subconscious hormonal excitement makes us overlook problems, like bugs. We call this **Cat Thinking**.

We fall for the same Cat Thinking trap when we enthusiastically buy into the latest digital promises and don't see the problems. It happens in digital healthcare all the time.

Digital healthcare, especially AI, is all new and wonderful — purrfect, in fact. It's sold to make us more efficient. Our hormones make sure we feel good about it.

Unfortunately, bad things like errors and harm do happen. There are plenty of stories about digital problems later in this booklet. But, because we're so happy with digital, we feel that any errors must be the nurses' or the doctors' fault. We are blind to any fault with the purrfect digital stuff.

Cat Thinking might sound trivial, but its consequences are very serious.

Maybe a patient dies from an overdose, and the investigators say the technology worked correctly, as designed. They therefore think the staff must have caused the problems because digital is wonderful. It purrs, and anyway it was very expensive and bought to solve such problems. Sack, discipline, or imprison frontline staff and "the problem's solved" — except it's a misunderstood problem, and it hasn't been solved. Poor system design was probably the key factor.

Most of us are eager consumers of exciting new things, and everything they promise. So we're culturally and hormonally driven to overlook how things go wrong because of poor design and bugs.

We must start thinking much more carefully and critically in healthcare. Improving healthcare isn't just a matter of getting more and new exciting digital systems, we need to transform the culture of digital healthcare. "New" won't fix anything, unless the culture that created and sustained the old problems improves.

Fix IT page 25

Don't get sucked in by new technology. Be curious and think critically.

IN THE BEGINNING

Back in 2000, I was a professor of computer science, teaching and researching how to design computers to be safe and easy to use. Then Nick, one of my students, was run over and ended up in hospital.

I went to visit Nick, and found him on an infusion pump with a Post-It note stuck on it. It said "Don't press this button!" I wondered why not, and as I did research in this sort of thing, I bought an identical pump and studied it.

I was surprised how bad it seemed, but maybe I didn't understand how infusion pumps should be safely and properly used.

So I found an anæsthetist who was interested in safety, and he allowed me to "go under cover" and become a porter for a week so I could watch operations that used the pump.

In every one of the six operations I attended, something digital failed.

The simplest story is when the ventilator crashed without any warning. It just stopped working. It was literally the blue screen of death from Microsoft Windows.

The anæsthetist had to reboot it and re-enter all the patient details all over again, so the patient wouldn't suffocate.

Fix IT pages 35 & 38

Since that experience with infusion pumps and ventilators I have continued to research digital health, to see how it can be made easier to use and safer. My research culminated in my book, *Fix IT: See and Solve the Problems of Digital Healthcare*.

Fix IT is published by Oxford University Press, and rather late in the day they got in contact with me about a missing picture. They were worried that there was a page with a caption but no picture to go with it. There was just a large blank gap. What did I want to do?

I explained. On one page there is a box listing all the things an anæsthetist needs to know to qualify. Once an anæsthetist passes they can, for instance, use a ventilator.

On the opposite, facing, page there's a box containing everything a programmer is required to know before they can develop ventilators.

It's completely empty — which was intentional — as there is no requirement that medical device programmers are qualified in any way or must have studied anything relevant. Neither the patient nor the anæsthetist has any reason to trust the ventilator. All the anæsthetist's skill in pressing the right button at the right time can be undone by its poor programming making it do something unexpected.

Fix IT page 39

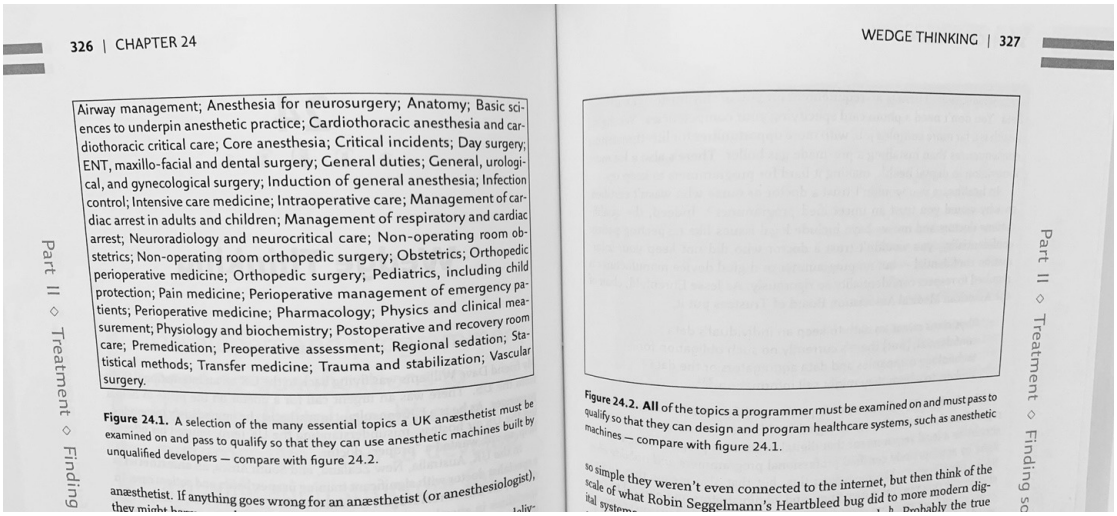

326 | CHAPTER 24

Airway management; Anesthesia for neurosurgery; Anatomy; Basic sciences to underpin anesthetic practice; Cardiothoracic anesthesia and cardiothoracic critical care; Core anesthesia; Critical incidents; Day surgery; ENT, maxillo-facial and dental surgery; General duties; General, urological, and gynecological surgery; Induction of general anesthesia; Infection control; Intensive care medicine; Intraoperative care; Management of cardiac arrest in adults and children; Management of respiratory and cardiac arrest; Neuroradiology and neurocritical care; Non-operating room obstetrics; Non-operating room orthopedic surgery; Obstetrics; Orthopedic perioperative medicine; Orthopedic surgery; Pediatrics, including child protection; Pain medicine; Perioperative management of emergency patients; Perioperative medicine; Pharmacology; Physics and clinical measurement; Physiology and biochemistry; Postoperative and recovery room care; Premedication; Preoperative assessment; Regional sedation; Statistical methods; Transfer medicine; Trauma and stabilization; Vascular surgery.

Figure 24.1. A selection of the many essential topics a UK anæsthetist must be examined on and pass to qualify so that they can use anesthetic machines built by unqualified developers — compare with figure 24.2.

anæsthetist. If anything goes wrong for an anæsthetist (or anesthesiologist),

Part II ◇ Treatment ◇ Finding

WEDGE THINKING | 327

Figure 24.2. All of the topics a programmer must be examined on and must pass to qualify so that they can design and program healthcare systems, such as anesthetic machines — compare with figure 24.1.

so simple they weren't even connected to the internet, but then think of the scale of what Robin Seggelmann's Heartbleed bug did to more modern digital

Part II ◇ Treatment ◇ Finding

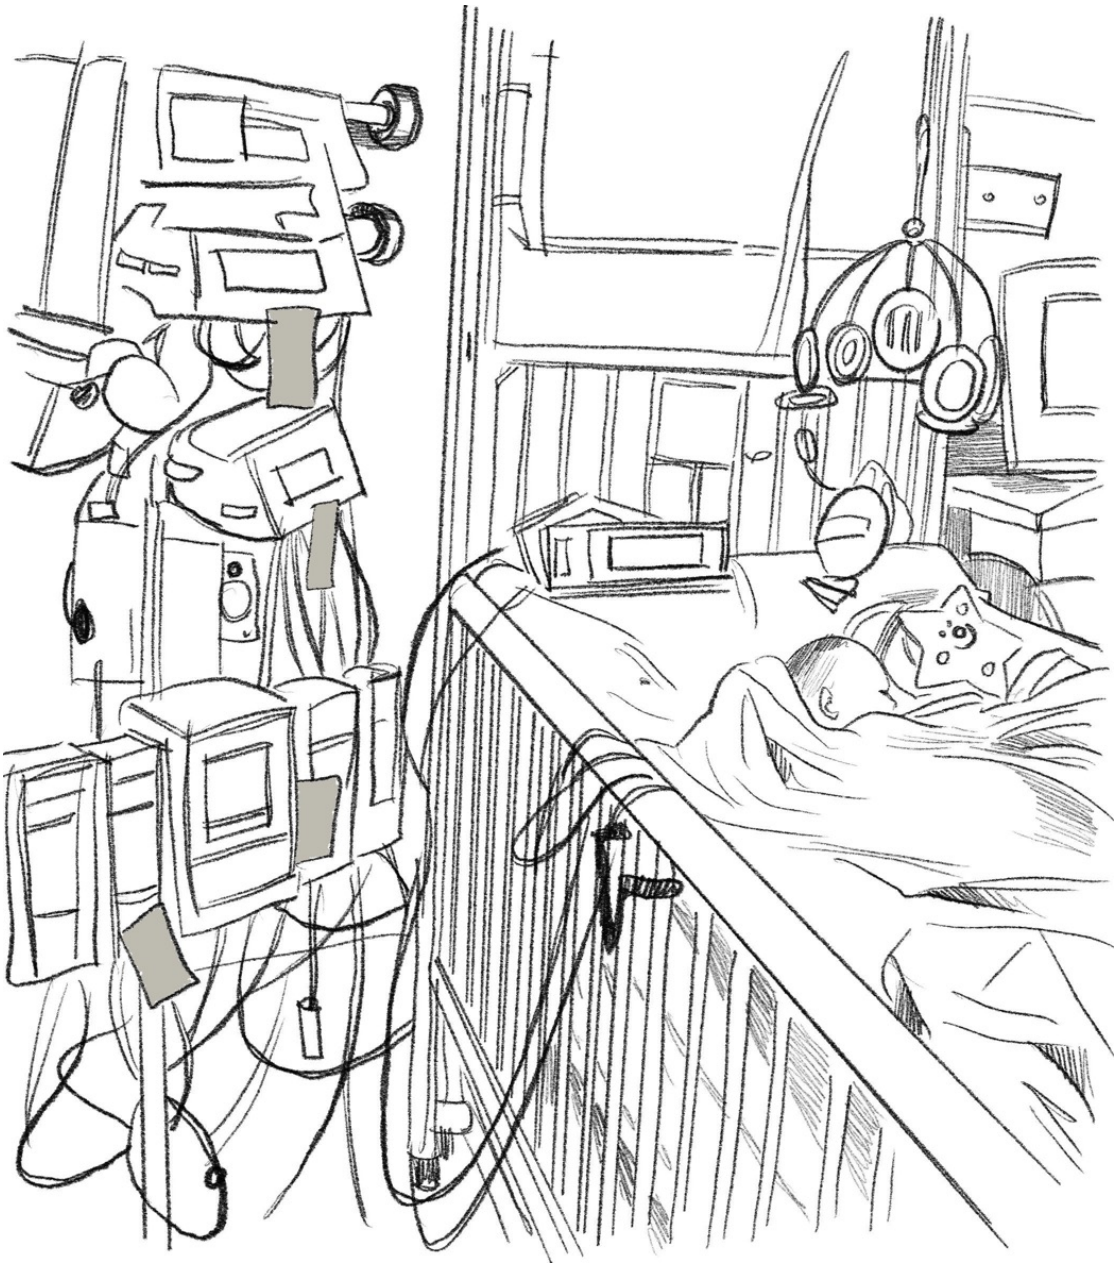

Fix IT page 39

If you want to be an electrician and wire up a hospital, the law says you must be qualified and up to date with the regulations. Yet if you want to build medical devices, ventilators, diabetic glucose meters, medical apps, X ray machines, GP appointment systems — anything digital in healthcare — then you don't need any qualifications at all.

That's why the ventilator in the story crashed. It was designed and developed by people who were out of their depth. It's impossible to assess people's programming skills without qualifications. Nobody is required to have digital qualifications for IT management either, so the ventilator was bought and used by people who had no idea it was poorly programmed. It should have been a lot better.

People who develop for digital healthcare need appropriate software engineering qualifications. This should be required by law.

BUGS AND MORE BUGS

“My heart leapt with joy! We had been trying for a baby for several months and now my period was three weeks late and the test was positive. The wave of joy was followed by a twinge of anxiety which came from having watched my brother grow up with all the challenges of having Down syndrome. We had decided to have the Down's test and then we could plan what to do.

Fortunately, the test was all clear, and I began to relax and look forward to meeting our baby.

Everything was going to plan, the baby's room was all ready, the excitement was building. I'd less than a month to go. Then, and I will never forget that moment, I read in the local paper that the software for testing for Down syndrome in our area had errors in it, and many women's results were inaccurate.

The floodgates of emotion opened, and when my partner came home from work, I was a blubbering wreck on the floor. Suddenly I had no idea what the rest of my life would be like . . .

Fix IT page 33

Right from January, midwives worried about getting unusual Down screening results back. Something was wrong. In April a midwife coordinator rang the Immunology Department, and spoke to Mr M. She wasn't just querying one odd result, but months of odd patterns of results.

Mr M believed it could wait for Mr K, who was on leave. Despite the standard incident reporting procedures, Mr M did not record the midwife's phone call.

Over a month later, after more phone calls, letters, and faxes expressing concern, Mr K finally realised there was indeed a problem. He found that 7,000 patients might have been affected.

Apart from a problem gradually escalating while nobody took it seriously, what was the cause of the problem?

The Millennium Bug is what had gone wrong, and the programmers and other technicians hadn't realised. They thought all the individual problems were unrelated, and that they were not significant. Besides, they had had a Millennium Bug workplan, and they didn't expect any computer problems.

Yet the Millennium Bug was messing up thousands of patient ages, which were an essential part of the Down Syndrome screening calculations.

Fix IT page 500

Of the 7,000 mothers potentially affected by incorrect calculations, Mr K told the Inquiry that he had identified 150 calculations that had moved from low chance to high chance. Two terminations are known to have been carried out because of incorrect screening reports, and four babies with Down Syndrome were born to mothers who thought their tests put them in the low-chance group. A hotline was set up, run by midwives, to support concerned mothers.

BUGS AND MORE BUGS

? What were the problems? Programmers took short cuts that turned out not to work, resulting in a computer system having bugs that caused chaos. The problem was initially not taken seriously by technicians and clinical staff. Later, a major incident was reported, and an Inquiry was set up to find out what had gone wrong, and to report on changes that must be made.

Interestingly, the Inquiry report relied on advice from a clinician, who proposed — without realising — a buggy solution to the problem. Furthermore, the clinician knew and admitted that his suggested program did not handle leap years correctly, despite the year 2000 itself being a leap year. The Inquiry's new, recommended code was thus as badly designed as the original, and arguably worse because it came with the Inquiry's authority as a solution.

Here was a hospital developing, implementing, and investigating its own computer systems from scratch without any professional software engineering oversight.

The Millennium Bug. A computer needs to know the age of the patient when Down screening is done, as it uses their age to help calculate the significance of the laboratory test result.

I was born in 1955. Throughout the 20th century, I could work out my age easily. For instance, in 1980, I was 80–55 or 25 years old, give or take a few months.

In the new millennium, the computer program's childishly simple age calculations all went wrong. My age in 2000 would have been worked out as 00–55, which would have meant I was apparently **minus** 55.

This bug affected all the Down Syndrome screening results.

Not only did the Down program have bugs, which messed up patient ages, but it didn't do a "sanity check" that the patient ages it was using made some sense (like being numbers between 10 and 100), so it went ahead and calculated bad test results from crazy ages calculated by bad code.

Any critical system should be programmed to report and block wild results being used. This one didn't.

In addition, the patient's age was not reported on any screening results sent to the midwives, so the midwives were not able to see when incorrect ages had been used.

In short, there was no professional error checking anywhere.

The Millennium Bug was an international failure. Millions of developers had made the same elementary mistake: hardly anyone had ensured their programs would work correctly after 1999. It might feel like ancient history, but everything in the Millennium Bug failure continues to be typical of problems with today's digital healthcare.

In 2011 the NHS's National Programme for IT (NPfIT) became the world's largest single IT failure. NPfIT was cancelled after wasting around £20 billion, and nothing useful was learned.

In 2018 a bug in a GP system resulted in 150,000 patients being affected by a data breach.

In 2019 an NHS clinic shared the personal emails of 37 patients using HIV services.

In 2020 nearly 16,000 COVID-19 cases were lost due to naïve Excel programming, and around 50,000 people were lost to contact tracers.

Almost every NHS system *still* has problems of being slow, needing another password, not being interoperable, needing upgrades, ignoring critical errors, and worse. Politicians continue to set ambitious goals with seductive promises for digital health, without prioritising dependability, safety, reliability, interoperability, usability, workload, maintainability, cybersecurity, etc.

The small selection above represents persistent digital incidents. The common underlying causes are an international problem.

Fix IT page 19

You depend on digital. Are the systems you rely on dependable? Were their developers qualified to build safe systems?

A HOSPITAL-WIDE PROBLEM

“I was standing in my kitchen and checking my work emails before I cooked the family dinner. That is when the bombshell hit. The email said I was not to go into work the next day. I was suspended.

I sat down at the kitchen table trying to take in what was happening. Trying to make my eyes focus and read the detail of what the email was saying.

I loved my job as a nurse on the cardiac ward. I loved the banter with the patients. How could something have gone so wrong as to lead to suspension? The email outlined that I was being accused of falsifying glucometer readings. We always had a number of diabetic patients on the ward, and we recorded their blood sugars regularly. I couldn't believe that I had somehow recorded them wrongly.

I rang my friend who I worked with. I was embarrassed, but I needed to talk to someone. She had just had the same email!

It turned out that about 70 nurses had been suspended. That helped — to not feel I was on my own. To not feel like the one black sheep that had let my patients down.

I thought long and hard about how I took and recorded the readings. I really couldn't remember ever making up numbers or forgetting to record them. My brain went over and over it, and when they said things would be easier if I pleaded guilty and admitted my errors I really didn't know what to do. If I pleaded innocent, but was found guilty, I would get a harsher sentence, maybe even go to prison. They said there was lots of computer evidence that showed I was guilty.

How could I argue against that?

Fix IT page 92

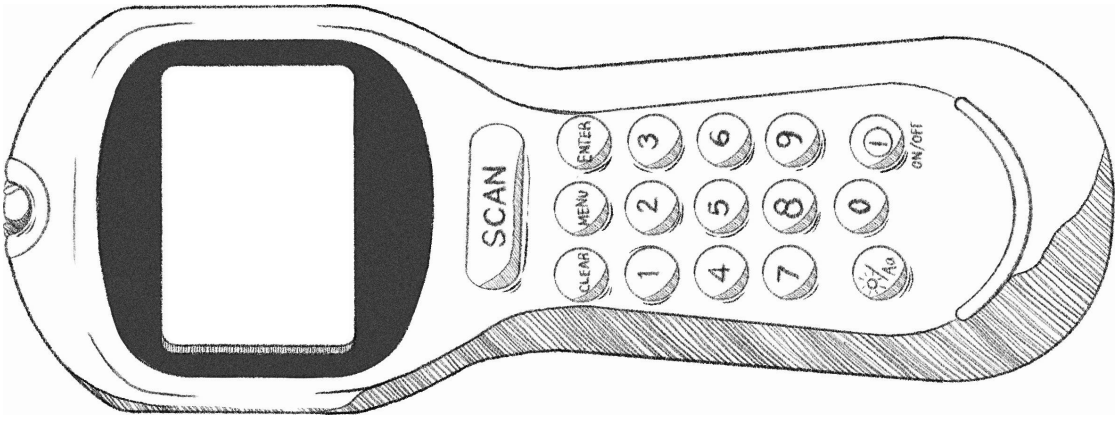

A HOSPITAL-WIDE PROBLEM

I was asked to be an expert witness in this court case involving two nurses who were being prosecuted for criminal negligence in their care of diabetic patients. The case was part of a large incident in the Princess of Wales Hospital where 73 nurses had been disciplined because computer records showed that they'd failed to do their job professionally. The police prosecuted some of the allegedly worst offenders.

Before a criminal trial, the defence and prosecution meet to review the case and see if it really needs to go to court. I was baffled how 73 nurses could all be so bad in exactly the same way, so I asked how the prosecution could believe all 73 nurses had made the same mistakes, when there are far simpler explanations like computer error, cyberattacks, or even an IT person with a grudge? The prosecution replied that all the nurses were in it together.

The case went to court, with two nurses in the dock. I was cross-examined almost every day, and ridiculed by the prosecution. Who was I to claim that a system made by a major international company, Abbott, was unreliable?

One day I said something that was too technical for the prosecution, who decided to call Abbott's Chief Engineer to respond to my technical points. When the Chief Engineer testified, he happened to say he'd visited the hospital. I prodded the barrister in front of me, and told her to ask "what did he do?" The Chief Engineer said a bit more about what he'd done. I asked the barrister, "when did he do that?"

I told the barrister: that's exactly when a lot of data vanished. Soon the judge intervened, himself examining the Chief Engineer. He then said the court would be adjourned while he wrote a ruling.

The judge read his ruling to the court. He said the prosecution had wasted everyone's time, and benefitted nobody. He ruled that the computer evidence was of no value, and asked the prosecution what they wanted to do . . .

The prosecution admitted they no longer had a case. The judge called the jury in, and told them there was no case to answer. The jury foreman rose to tell the court that there was no case to answer.

The judge said, "Release the prisoners."

? What were the problems?

The Princess of Wales Hospital had problems with their computer records. They called in an engineer to sort out the problems. It became clear in court, the engineer had *really* sorted the problems out, by deleting the problematic data.

Later the hospital relied on the "sorted out" computer systems to provide evidence that nurses had failed in their standard of care. It didn't cross anyone's mind that the computer evidence was unreliable, even as they disciplined one nurse after another.

The scale of the problem became so large the police got involved, and it became a criminal investigation. The mindset that the nurses were to blame became entrenched, and because there was a criminal investigation nobody would talk openly about anything.

It's worrying the hospital did not detect lots of data being corrupted. How would it have noticed had a cyberattack corrupted data?

It's even more worrying when you realise that the same and similar systems will be in use in thousands of hospitals around the world with similar lack of monitoring.

Once a case like this reaches the criminal court, a "legal presumption" in the UK makes it very hard for people to defend themselves. In law, the presumption is that computer evidence is correct.

It thus becomes almost impossible for defendants to challenge computer evidence because, since the evidence is presumed to be correct, how can you find out any reasons why the program code or documentation should be disclosed for your team to review?

Be slow to blame staff. Ask — could computers, devices, and records be unreliable?

A COUNTRY-WIDE SCANDAL

Temporary picture

Seema Misra became a post office operator in West Byfleet, Surrey, in 2005. Over the next two years she had attempted to balance her books, borrowing money and transferring takings using the Horizon accounting system provided by the Post Office.

Seema regularly reported her problems to the Post Office Helpline — several times a week, and each time she was told to "roll over" the accounts or balance them with her own cash.

Finally, she failed to keep her head above water.

After an audit of her accounts found a discrepancy of £74,000, she refused to plead guilty and was prosecuted by the Post Office.

During the trial, she found she was pregnant. It came as a complete shock when she was convicted of theft and false accounting. Seema collapsed in the dock when the judge read her sentence, sending her to jail. She was ordered to pay compensation to the Post Office.

The media called Seema a "pregnant thief." Her husband, Davinder, was beaten up by locals, who accused them of coming to the UK "to steal old people's money."

In another world, Jarnail Singh, the Post Office's senior criminal lawyer, celebrated Seema's conviction. He sent an email to Post Office executives writing "It is hoped the case will set a marker to dissuade other defendants from jumping on the Horizon bashing bandwagon." He added: "Through the hard work of everyone . . . we were able to destroy, to the criminal standard of proof, every suggestion made by the defence."

In prison, Seema was put on suicide watch. She says, "If I hadn't been pregnant I would have definitely killed myself."

In fact, Seema's £74,600 shortfall was caused by bugs in the Horizon computer system.

It was not until 2021 that Seema's conviction, along with 38 other postmasters' convictions, was overturned by the Court of Appeal. The Court of Appeal ruled that her prosecution by the Post Office was abusive, an affront to justice, and an affront to the conscience of the Court, a rare and extraordinarily severe ruling. It had taken eleven years to prove Seema's innocence.

Seema is one of more than 700 post office operators wrongly prosecuted for fraud, theft, and false accounting.

Fix IT page 508

A COUNTRY-WIDE SCANDAL

The Post Office Horizon scandal is basically the same story as the nurses suspended over alleged negligence. The Horizon scandal is in this booklet because it shows how serious digital problems are widespread, widely ignored, and ongoing.

An organisation gets an expensive digital system, which promises to revolutionise how work will be done. In fact, both Fujitsu (for the Post Office's Horizon system) and Abbott (for hospital blood glucometers) claimed their systems would reduce error.

Then someone is found to have made some mistake. Stolen money or not performed a test on a patient. Such things happen, regrettably. There are procedures to follow.

A few days or weeks later, another guilty person is identified. And, again, procedures are followed. And so on.

Soon the organisation realises its reputation is at stake. The media gets interested. The organisation knuckles down as the number of "guilty" people rises.

Nobody wants to stop and think that their computer systems may be unreliable. In the courts, the computer evidence is presented as infallible.

The defendants are told the evidence against them is overwhelming. "Do you remember *exactly* what you did several years ago? Well, the computer does."

So you are sure to be convicted, and you should plead guilty to reduce your sentence.

Then, in court, the law backs up the prosecuting organisation. It is presumed that the computer evidence is correct. This presumption means that the prosecuting organisation does not need to justify to the Court why it thinks its evidence is correct.

In the Post Office case, it's clear that at some point the Post Office knew the evidence it was using was wrong, and it then went into a major cover up.

The Post Office relied on a detail of British Common Law: computer evidence is presumed correct.

The legal presumption was recommended by the Law Commission. There's good evidence that they, the people at the top, also misunderstand computers.

There are huge differences in the Post Office and Princess of Wales cases.

The Post Office Horizon IT Inquiry is finding many examples of the Post Office's willful misrepresentation of facts about Horizon. The Court of Appeal ruled that the Post Office's convictions were achieved by a serious abuse of the Court system. The Post Office prevented a fair trial, and the Court of Appeal called the Post Office abusive.

In contrast, senior staff at the Princess of Wales Hospital believed the computer evidence and never realised otherwise. It seems surprising, though, that they hadn't the curiosity to investigate why so many nurses apparently failed in exactly the same way.

The Post Office Horizon IT Inquiry has been shown some of the computer code inside Fujitsu's Horizon system. Here's a short example.

The idea here is to make the value of d negative. You would imagine the code would say something like "set d to $-d$."

Instead, it says "if $d < 0$ then set d to the absolute value of d, else set d to d minus twice d." If you wanted to make something very straight forward very obscure and error-prone, hard to understand, and hard to debug — this is a good way to do so. It's coding to be ashamed of, and has no place in professional code.

Professional programmers will notice that this peculiar code is not *exactly* equivalent to set "d to $-d$," as it introduces rounding errors, which will make things worse in subtle ways.

Lots of Horizon code has other bizarre problems, but the bugs are impossible to explain briefly as the code is so confused.

The legal presumption that computers are reliable, and similar mistaken assumptions in internal investigations, need correcting.

HEADS IN THE SAND

“As I pushed open the door of the hospital, I could immediately see that something was wrong. There were people milling around the foyer looking like they didn't know where to go. The volunteers on the reception desk were on the phone. One was writing out a big notice. I saw the word “cyberattack.”

It was 14 May 2021, and I was working as a doctor in Dublin, Ireland. We were already in the middle of a pandemic, and a cyberattack would make things go from bad to worse.

And it did go from bad to worse.

The chaos that followed was indescribable. There were no lists of patients for my clinic. I couldn't access any medical records online. I couldn't print labels for blood tests. One of my patients needed a scan to see if his cancer had shrunk. I had to go down to the scan department to view the scan because they couldn't send it up to me. But then I couldn't tell the patient whether he was getting better because I couldn't access his last scan to compare it.

All computer screens were down — nothing was working.

We had to shut the radiotherapy unit whist we worked out how to safely calculate and administer treatment. Patients kept arriving as we had no way of cancelling appointments because we couldn't even access the booking lists and find out who had appointments.

That night we had an emergency meeting to pool ideas on how to keep operating safely. We started using WhatsApp to communicate and send images. The criminals demanded $20 million. They said they already had 700Gb of our patient data that they would be releasing.

The stress of caring for people and keeping them safe has been indescribable.

Fix IT page 214

The HSE, the Health Services Executive, is basically the NHS of Ireland. It has over 130,000 staff who are dependent on connected and reliable digital systems.

The Conti cyberattack started on 16 March 2021. On 18 March, somebody opened a phishing email, which included a malicious Microsoft Excel spreadsheet. The user's access rights let the Conti cyberattack in. The cyberattack at first spread silently.

There were in fact several earlier occasions when people noticed the beginning of the attack, but these were not declared a cyber incident, so opportunities to start fixing the problem before it had seriously escalated were missed.

The full cyber onslaught detonated on 14 May. Because Conti encrypted data, staff across Ireland lost use of every IT system, including patient information systems, clinical care systems, laboratory systems, financial systems, payroll, and procurement systems. Even monitoring sterilisation. Everything went down.

Email and even phone lines went down too. There was no way to communicate with HSE's national centre. Everyone had to revert to pen and paper to work and continue patient care, or improvise using their own phones with apps like WhatsApp.

The chaos lasted over six months.

Given Conti happened during the pandemic, there was a terrible load on staff. The cyberattack management team formally brought in Occupational Health because staff mental health had become a critical issue in managing the recovery. Clinical Indemnity had to be provided to doctors, nurses and midwives, because there were no patient records and nothing worked.

It took till the end of September 2021 to sort it out, when all servers were considered decrypted, and with most, but not all, applications restored back to use.

In the UK, the Health and Social Care Act requires compliance with standards, like DCB 0129 and DCB 0160. These standards require digital risk assessment registers, and more, but don't require competent, appropriately qualified external oversight. So if an organisation doesn't fully understand digital — which is routine — it may easily think it complies with standards, when actually it doesn't understand the technicalities and fails to comply with them. Worse, the standards themselves are weak, because, apparently, their authors didn't know how to write rigorous requirements for dependable software.

HEADS IN THE SAND

The gang who developed the Conti cyberattack used off-the-shelf tools, readily available to any hacker, to create it. It was basic stuff. It could have been much, much worse had the hackers had any serious plan. Their attack didn't spread from HSE's internal systems to external cloud systems, for instance.

The hackers had encrypted data, so making it useless. Maybe because they realised their attack had had a terrible effect on an entire nation and they were out of their depth, they soon released a decryption key. It wasn't perfect but made recovering data a little easier. Malicious hackers could have deleted data (and data on backups), or, worse, corrupted it so that patient records were wrong. Malicious hackers could have targeted patients, famous patients, or types of patients the hackers didn't like. They could have changed blood results or drug regimens — or anything.

Fortunately, the Conti attack was technically trivial stuff.

HSE, like everyone else, had already had experience of cyberattacks like WannaCry. So the question arises why had HSE not kept up maintaining and developing defences against cyberattack?

The Conti attack exposed that the Irish national health service was operating on fragile IT with a system that had evolved without much thought, rather than been designed for resilience and security.

It wasn't just the computer systems, but the management systems too: management had no professional cybersecurity expertise. The report into the Conti attack says "The HSE also had only circa 15 full-time equivalent ('FTE') staff in cybersecurity roles, and they did not possess the expertise and experience to perform the tasks expected of them."

These people were in principle responsible for supporting 130,000 workers and giving specialist guidance to everyone how to recover — when they had no obvious means to communicate with them. A tall order for inexperienced people during a disaster.

? Have you had cyber-training? Some cyber training tries to make up for bad digital system design. Ironically, in Ireland, the people who did best during the cyber-attack were those who still relied on paper instead of digital. Going paperless is often held up as an ideal, but it has many drawbacks, especially when it is not designed or managed well.

You can explain a cyberattack away by saying it was caused by criminals, or possibly by under-trained staff who had not done enough cyber-awareness training. All true. A more strategic understanding is that your systems are inadequate. They had bugs and backdoors that should never have been there. The manufacturers should have maintained the systems. Updates the manufacturers distributed, if not automatically applied, should have been routinely applied by local IT staff.

There should be no legacy (out-dated) systems in healthcare.

Manufacturers should have warned and helped transition to more modern, safer and more secure systems — after all, they built in the original weaknesses that the criminals exploited.

Manufacturers' warranties and contracts typically exclude digitally-related liabilities. On the contrary, we should contractually require manufacturers to guarantee that their systems are, as far as reasonably practical, free of safety-related errors throughout their lifecycle. One way to help this is by committing to continual iterative design, and regular maintenance.

If a manufacturer won't stake their business on the quality of their products, don't stake patients' lives on their products either!

Fix IT pages 298 & 313

Check the cybersecurity compliance of systems you rely on. Frontline staff should not have to learn how to compensate for poor system design.

AS IF BY MAGIC

RaDonda Vaught was familiar with the BD Pyxis, an automatic drug dispensing cabinet, at Vanderbilt University Medical Center (Tennessee, USA). In December 2017, RaDonda needed some Versed, a sedative, to help an anxious patient, 75 year old Charlene Murphy, relax and have her MRI scan without worrying.

To get Versed out of the drug cabinet, you would type VE, and select Versed from the screen showing everything starting with VE. Once a drug is selected, a drawer with several boxes will open, and the right box will pop open automatically, and you take the drug out.

But this particular day things didn't work like that.

RaDonda couldn't find any Versed. As was routine in the hospital, RaDonda did an "override" to get a larger list of drugs, to show more options starting with VE. It was common practice to override the cabinets, as they were so hard to use. Some reports say there was a "persistent software problem."

RaDonda entered VE again, a drawer opened, and a box popped open.

RaDonda took the drug out and gave it to Charlene. Unfortunately, the drug wasn't Versed, and Charlene died soon afterwards because she'd tragically been given Vecuronium, which is a paralytic.

In fact, the Versed was stored in the cabinet only under its generic name, Midazolam, so typing VE would never have found it.

Interestingly, the hospital had reprogrammed its drug cabinets to accept short 2 letter abbreviations to make them easier to use, despite 5 letters being widely recommended for safety reasons. Note that the first 5 letters of Versed and Vecuronium are different. Following best practice would have blocked RaDonda's error.

RaDonda reported the error. She was fired from the hospital, arrested and prosecuted. She was found guilty of gross neglect of an impaired adult and negligent homicide.

Fix IT page 172

The case triggered a widespread outcry. The American Bar Association says, "A robust culture of safety relies on self-reporting and transparency to drive process improvement, and criminalising errors instead foments blame and creates fear."

Paul Curzon tried out a magic card trick on me. I had no idea how he kept identifying the cards I had secretly picked. He kept on doing it! Everyone watching was laughing, but I just couldn't see what he was doing to keep catching me out.

Paul was deliberately, and very successfully, fooling me by carefully controlling my attention, memory, and expectations.

Magic like this proves how easy it is to manipulate things so that people will predictably make the same mistakes. Magicians do it for entertainment.

Conmen do it to cheat you. Email phishing tricks people out of their money — or an entire country out of having working healthcare.

Digital healthcare tricks us by accident, but unfortunately accidental misdirection is just as good at tricking us as deliberate magic tricks.

The psychology of predictable error, which is used to trick us in magic, can be turned right around into the positive science of Human Factors: how to manage and avoid error, how to avoid being tricked into unwanted surprises. Crucially, Human Factors can help designers avoid building misleading digital systems.

Medical systems used in hospitals accidentally perform "tricks" on staff every day. Staff miss critical details because they are misdirected by a toxic mix of clinical pressure and poor design.

Fortunately, healthcare staff are professional, and usually they very soon notice problems. But sometimes, they have no idea they have been tricked, and the consequences may be catastrophic for patients.

Like me with Paul's card tricks, staff trying to use a complex device — like an automated drug cabinet — are surprised and confused about what's happened. That makes them very easy to blame when mistakes happen.

When digital systems mislead us, it's because negligent design hasn't taken proper account of Human Factors. Incident investigators, too, are often unaware how common poor digital design is the "root cause" of error.

Fix IT page 262

AS IF BY MAGIC

Temporary picture

Paul Curzon tricking an audience. Pick a card . . . any card

“I am the head of pharmacy at a large American hospital. Digital medicine and prescribing is a really exciting area for innovation and improving efficiency. Our hospital is at the cutting edge.

I've been in lots of meetings discussing AI and robots. We had a campaign to get approval to buy the latest technology. Last month, we were one of the first hospitals to install a $7 million pharmacy robot.

The new robot promises to eliminate human error. But only today, a doctor entered a dose of 160 milligrams per kilogram of patient weight, when it should have been 160 milligrams full stop. So the intended dose of 160 milligrams was multiplied up by the patient's weight. This was a very large overdose.

The wrong dose wasn't noticed by the system, but was packaged up by the robot, and the patient was given this dose by a nurse who was just doing what the robot told her to do.

Very sadly, six hours later, the patient fitted and stopped breathing.

This has brought us down with a huge bump — it's tragic that $7 million of the latest AI and robots isn't sufficient to eliminate error.

Fix IT page 126

? Why wasn't such a dangerously wrong dose noticed by the hospital's new robot? In fact, clinicians had been overwhelmed with irritating computer error messages interrupting them all the time, so the hospital medical center had decided to disable many alerts. Of roughly 350,000 medication orders a month, there were now “only” 17,000 alerts a month. Unfortunately, the mistake described here was not spotted because it was one of the blocked alerts.

Understanding how design error, Human Factors, and the psychology of error combine with poor digital design is essential for incident investigators.

In 2021 a strong candidate applying for a training post in anæsthetics in the UK was rejected by the digital recruitment system, Oriel. The candidate was surprised to find out they were "unappointable." After an investigation into the complaint, it was found that 34 more candidates had been affected as well.

Separate Excel spreadsheets to record candidates were created by people across the country. The spreadsheets varied depending on who had created them, which made them hard — and error-prone — to combine into a final master spreadsheet. The inquiry into the incident also found that few people developing the spreadsheets understood the behaviour of Excel's VLOOKUP function, which had been used in the spreadsheets.

Fortunately, nobody was harmed by this interoperability fiasco, but the story is a warning for anyone developing digital systems, including spreadsheets, that will be used for clinical purposes.

? What was the problem? People drift into creating applications that require professional software engineering skills to do well. Because they aren't software engineers, they don't notice the transition from what seem easy and obvious ideas to needing professional skills as things scale up. They likely don't understand the very basic software ideas, like specifications, assertions, invariants, testing, pair programming and code review, or why these principles are essential to develop systems reliably.

The story above shows how Excel, a powerful and flexible system, was too flexible and complex for the people using it. Its power created problems they didn't recognise or know how to avoid. They never got professional programming help, which clearly they really needed.

The same trap swallows up developers too. It's very easy to program and "make things appear to work," but it's *very* hard to make things work safely under all circumstances.

A common way to cause an unnoticed error is, ironically, to correct an error that you *do* notice. Many devices don't work like familiar systems, such as your phone and PC.

Imagine you are to enter 5.5 mg as a drug dose on an infusion pump. Unfortunately, you accidentally press the decimal point twice in a row as maybe the key bounces. Fortunately, you notice your error: you know you typed 5.. by mistake. You press the delete or backspace key, which you know will correct the number back to 5., then you can press the 5 key, so in the end you'll have entered 5.5 as intended.

Unfortunately, the patient gets 55 mg not 5.5 mg, because almost all devices delete *both* the decimal points. Worse, if this bug results in an incident, the device's log will show what *it* did, not what *you* told it to do.

Other devices have other bugs. For example, entering 100.5 on some infusion pumps results in 1005, another out by ten error, because they silently ignore decimal points in large numbers. Of course if you now hit delete twice because you didn't mean to press that .5, you'll get 10 not 100 as you expected!

Fix IT page 177

According to the *Daily Mail*, Arsula Samson died because a "blundering nurse" entered a ten times overdose. No error was found with the infusion pump (!), so investigators ruled the death was due to "individual, human error." The coroner's verdict was accidental death to which neglect contributed.

The action plan after the death saw medical staff retrained, and new infusion pumps and software brought into all wards "to reduce the risk of error." If expensive training, new pumps, and software were needed then the system must have played the critical part in inducing the tragic error.

Fix IT page 71

The main reason why digital things go wrong is that people — whether clinicians, administrators, investigators, technicians or developers — don't notice when technicalities start going over their heads.

Know your limits. If you don't know your limits, and aren't aware of them, you aren't safe developing digital systems — even simple Excel spreadsheets.

WHAT TO DO IF THINGS GO WRONG

If you end up in a disciplinary hearing, or in litigation, here's what you should consider:

1. You are not alone. Talk to people. Try to find out if other people have had similar experiences — if so, suspect the system, not the people.

2. Check out the NHS Just Culture — for example, it says "Are there indications that other individuals from the same peer group, with comparable experience and qualifications, would behave in the same way in similar circumstances?" — if so, this would be a likely sign of digital problems!

   A clear guide to Just Culture is available at
   https://www.england.nhs.uk/patient-safety/a-just-culture-guide

3. It is important to talk to a competent expert. Although Expert Witnesses are available for court work, a good place to start is to get in contact with your local university Computer Science Department, and talk to someone — maybe a PhD student or Research Assistant — who is working in usability, safety, cybersecurity, or dependability.

4. The freely accessible database *MAUDE – Manufacturer and User Facility Device Experience* has details of millions of medical device issues, and may be relevant to your situation. Search online for "FDA Maude."

5. Require the disciplinary panel or your prosecutors to provide the documents listed below. Where these items won't or can't be provided, everyone should treat evidence with appropriate caution. Your friendly expert (see point 3 above) may need to modify or extend this list, depending on your circumstances.

   Note that in a court case, a judge can order disclosure of facts to an Expert Witness if the organisation is not prepared to provide the documents or access to the systems to determine the facts yourselves.

   (a) A copy of the actual computer-readable evidence (not a paper printout) that they are relying on — so you or experts can analyse it. It must be examined by experts before anyone agrees that the system has produced correct output, and what its significance is.

   (b) Copies of independent certification to the relevant international standards that the digital systems are reliable, and adequate for evidential purposes.

   (c) Evidence of the relevant serial numbers and software version numbers. Without knowing these, it is impossible to confirm how the device(s) will have behaved. It's also possible that old or obsolete software version numbers will indicate that the systems have not been properly maintained.

   (d) Evidence that the systems are suitable for the clinical purposes for which they are used. If they cannot provide this evidence, then they cannot logically rely on any presumption that the electronic evidence they are using is reliable, Common Law presumptions or otherwise.

   (e) Proportionate evidence of the forensic standards followed, to ensure the integrity of the evidence before and during the investigation. Has the device been used since, or reset? Were there system failures or cyberattacks? Does the data include timestamps? Note, for example, with spreadsheet and simple database data, rows and columns may be edited, duplicated, or deleted without leaving any record of tampering or accidental edits.

   (f) Independent confirmation that the data in fact refers to you or your actions. In many environments, for historical reasons, there may be multiple staff cards in circulation with your ID on them — so "your" ID may not refer to you.

The Further Reading (page 27) has additional helpful resources.

There are things you can do to prepare for any investigation.

LEARNING FROM THE PAST

Deaths and illness around childbirth had always been a fact of life until 1847, when Dr Ignaz Semmelweis noticed that his hospital wards had a higher rate than the nearby convent hospital. He set out to find out why. What was he doing wrong?

Semmelweis started collecting statistics. He noticed that his wards had fewer sick patients in the summer. But why? He realised that in winter his doctors attended autopsies to learn about anatomy, then went straight back to the wards. In summer they didn't do this so often.

Semmelweis speculated something was getting back to the wards from diseased bodies in the morgue. He instituted hand washing. His intervention soon reduced maternal death from around 20% to 2%.

Unfortunately, Semmelweis met a lot of resistance to his ideas. Doctors didn't like being told they might be the cause of illness. After all, Semmelweis had no real theory why his ideas worked.

A real explanation had to wait until Louis Pasteur developed germ theory, which at first only explained fermentation. The Scottish surgeon Joseph Lister then connected germ theory to putrefaction and disease. Lister realised that antiseptics would destroy germs causing disease, and his success as a surgeon soon became famous. Healthcare was persuaded.

Antiseptics prevent but don't cure disease. Effective cures had to wait for antibiotics, which Alexander Fleming discovered when he identified penicillin in 1928.

Antibiotics are now widely used to treat infections. But in turn, using antibiotics is creating new problems in this invisible world of bugs. Bacteria evolve, and become resistant to antibiotics. Some antibiotics are thus losing their power, which can be catastrophic for infected patients — it's a new global health threat. Guidelines are being developed so antibiotics are only used when they are effective and don't increase antibiotic resistance.

Fix IT page 15

Bugs that cause disease are invisible, and for thousands of years, people could only speculate about illness. Most people just accepted disease.

Eventually, Ignaz Semmelweis worked out a cause of disease, but he met considerable resistance from his colleagues.

Digital healthcare is a new intervention, affecting all areas of healthcare. Digital health has hidden bugs. These bugs impact primary care, secondary care, public health, and even systems used by patients.

Managing digital bugs doesn't require antiseptic procedures, it requires computational thinking. Unfortunately, just like theories of infection in the nineteenth century, computational thinking meets resistance.

Why would we have bought an expensive AI system if it wasn't going to be an effective system? Just like Semmelweis's colleagues, we resist being told we may be wrong. Moreover, as cat thinking shows, the excitement about digital itself makes it much harder to think about its limitations and weaknesses.

This leads to the inevitable stories of unexpected failure and misattributed blame in the stories throughout this booklet.

Just like early ignorance of bacteria meant that curing disease was almost impossible, ignorance of computational thinking makes it very hard to recognise, talk about, or to address the problems of digital bugs.

? What does **computational thinking** mean?

We haven't defined computational thinking in this booklet. We haven't defined antibiotic either. You take for granted that to understand antibiotics you need some medical training. It's the same with computational thinking: to understand computational thinking you need computer science and software engineering training.

Chapter 13 in *Fix IT* is all about computational thinking.

Fix IT page 151

There are parallels between antibiotic resistance and digital failure.

Just like the discovery of antibiotics was revolutionary, and antibiotics were first seen as a "magic cure," so, too, people promote digital as a cure-all solution to many of today's healthcare problems. But it's counter-productive to try to solve problems by getting more or newer digital systems — digital transformation — without stopping to ask: Will it be better? Will it be safe and effective? When is enough enough? We won't know unless we understand the underlying computer science.

LEARNING FROM THE PAST

Medical device and system manufacturers often complain about regulatory burden — the cost to them of complying with even the current regulations around digital healthcare.

I remember a manufacturer saying that if they followed regulations to the letter, their innovative systems would be so delayed they'd be obsolete by the time they reached the market. I wondered, if things are going to get obsolete so quickly, why would anybody want to spend any money on them?

We wouldn't generally go along with such arguments for drugs — we'd first want to know drugs are safe and mostly free of side effects before they are marketed.

Pharmaceutical companies rarely complain about regulatory burden. They employ highly qualified chemists and microbiologists to develop and test new drugs. They take regulation in their stride.

Perhaps most of the "regulatory burden" digital companies complain about is brought on by having staff who find competent computational thinking too hard?

Inexperienced developers use what's called a "happy path" to test their systems. They simply check that things work as they are supposed to work. They can now tell everyone that it works exactly as it's supposed to work — implying any problems are your fault. But the happy path means they didn't explore possible failures, that is, how things *aren't* supposed to work.

Competent developers instead test all the possible paths where their systems are and aren't used as they are supposed to be. This is extremely hard, as there are an overwhelming exponential number of ways of using things in unexpected ways. Developers who want to do thorough testing therefore use sophisticated computer tools.

See page 26 for more on how competent developers work.

Digital conceals complexity. Computational thinking is the right approach to manage complexity.

SAFETY FIRST

My Dad died from an over-infusion. Apparently, a nurse had left a drip running unattended, and Dad went into cardiac failure from too much fluid. It would probably have been better if Dad had been on a digital infusion pump that had been programmed to stop at a preset maximum dose.

Here's a picture of me having an infusion of rituximab, using just such a digital pump. It was a long process. My dose needed adjusting throughout the day.

I was glad that the type of infusion pump they used with me, an Alaris GP, has up △ and down ▽ keys to adjust the infusion rate. Our research shows that △/▽ keys are about twice as safe compared to numeric keys for entering drug doses.

Our research has also shown that you can halve the "out by ten" error rates of numeric keys by better design ("out by ten" errors often happen with the decimal point in the wrong place).

Put another way, if you are using a suboptimal design, the design itself will cause errors.

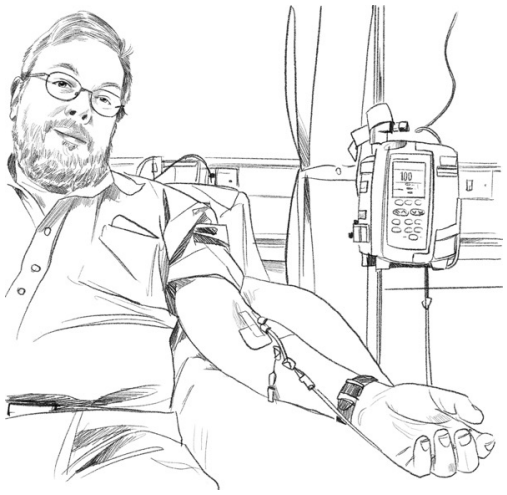

Fix IT page 412

? Digital devices vary enormously in quality, ease of use, safety, and in environmental impact. How does anyone know what a good robot or infusion pump is? How can hospitals choose and buy the best and safest equipment?

It's tempting to try to buy the cheapest. Running costs are a big factor. Every time an infusion pump is used, new lines are required — and these are often proprietary, only fitting the particular make of infusion pump. This causes "lock in" thus limiting choice. When a contract is set up to buy thousands of pumps over the next ten years, various deals will be done to balance and manage the short and long-term costs.

The procurement process rarely considers the long-term safety or usability, because nobody knows what anything's safety or usability is. They just assume these devices are safe and usable. Indeed, manufacturer's descriptions like "easy to use" or "eliminates error" mean nothing without rigorous evidence.

When we buy tyres for our cars, we have a very similar problem. We want the best tyres, but we don't want to pay over the odds. Some tyres are safer than others, but we have no idea which. Some are clearly cheap re-treads, whereas some are made by reputable manufacturers, sometimes with household names.

The European Union noticed this problem, and decided we, the customers, needed some help. Tyres now have rating labels, covering stopping distance, noise, and fuel economy.

When you buy a tyre you can now see how good it will be.

Fix IT page 401

Digital healthcare should be evidence-driven. More digital healthcare research should be done, and the research needs proactively adopting.

SAFETY FIRST

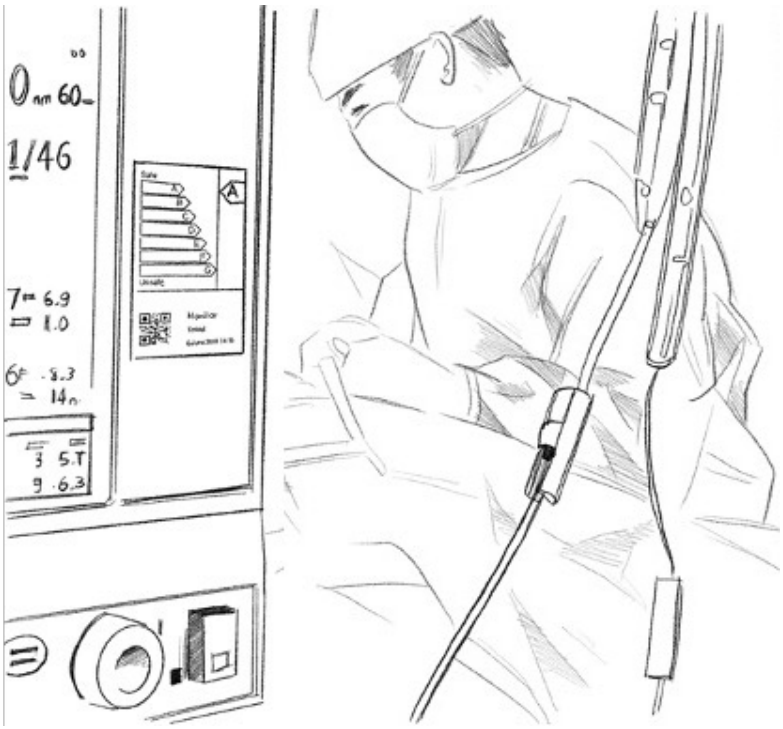

Safety labels would prominently show how safe and usable a device is. Staff and patients would all be more aware of how safety depends on quality equipment.

Ratings would combine several measures, such as the results of a bag of standard experiments with real and simulated users, just like car safety tests. The device illustrated in the picture is 'A' rated.

The safety label also has a QR code on it, which quickly provides relevant information about the device. The QR code would lead to a web page that provides access to the device's configuration, serial number and software version, and, ideally, its logs as well.

Fix IT page 405

There are similar styles of labels for rating the energy consumption of fridges, cookers, and washing machines. Not only do these labels help us know how to buy better products, but they have also improved the quality of products that manufacturers make. Manufacturers who made more efficient fridges sold more fridges. So manufacturers competed to make even more efficient fridges. Now fridges have energy ratings that are far better than the original A rating. In fact, the ratings have had to be adjusted to account for higher standards. Everybody benefits. It'd be nice if that sort of improvement happened in digital healthcare too.

We already require electrical safety labels (from Portable Appliance Testing, PAT testing) on everything plugged in at work, even though hardly anyone is harmed by electricity at work. So why not have digital safety labels too?

Fix IT page 402

We need safety ratings to inform procurement. Manufacturers will then make safer and safer systems as they try to out-compete each other once the quality of digital products is made obvious.

GETTING BETTER

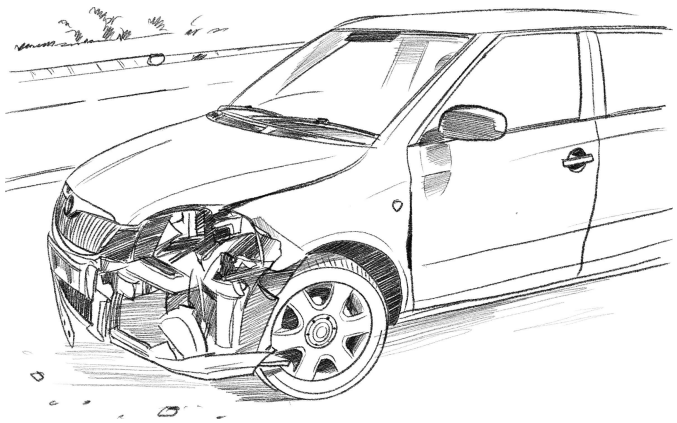

The day before he got married, our son Isaac asked to borrow our car, a silver Škoda Fabia. Well, of course he could!

About an hour later, the police telephoned to say there'd been an accident.

Isaac had collided with another car. The other car was on the wrong side of the road, and flipped over, ending up in a ditch.

You can see our car after the accident above. It has a crumple zone at the front, which has crumpled. The air bag also went off, and it saved Isaac from injury. This is what crumple zones and air bags do: they absorb energy in crashes, and save people.

There are many safety features in modern cars that are harder to see, like the seat belts, the ABS brakes, the "crash box" — the rigid frame to protect the passengers — and more. Manufacturers like Škoda are keen to promote their cars on the quality of their safety features.

Car manufacturers today want drivers to survive car accidents, or, better, to *avoid* accidents. If your tyres and brakes are good (and properly maintained) you can stop quickly in a controlled way, and you will never have an accident at all if your car stops before it hits anything or anyone.

Despite the speed of Isaac's crash, and the damage to the two cars, Isaac and the people in the other car walked away uninjured.

Safety technology works.

The deeper point is that errors will happen, but they need not lead to harm.

Fix IT page 127

? If Isaac had been driving a 1960s car, at the speed the cars hit each other, he'd likely have died, like in the awful picture below.
What changed to make cars safer?
What can digital healthcare can learn from industries, like cars and aviation?

NCAP, the New Car Assessment Program, started testing car safety in 1979, and has been instrumental in making cars safer. Car manufacturers now appreciate NCAP because they can promote how safe their cars are. People who buy or sell cars want NCAP ratings because they help them negotiate and think clearly about safety and quality. NCAP is so successful, it's gone global — see www.globalncap.org

What can we learn from NCAP's success to help improve digital healthcare quality?

NCAP has a mission statement, and by changing NCAP's "safer car" focus into "safer digital healthcare" we get a powerful mission statement for digital healthcare:

1. Support safer digital healthcare in emerging markets by offering support, guidance, and quality assurance.
2. Provide a co-operation platform for healthcare organisations around the world to share best practice, to further exchange information, and to promote the use of information to encourage the manufacture of safer digital healthcare across the global market.
3. Promote digital healthcare safety with proven effectiveness, and encourage its accelerated use across the globe by increasing awareness and, where appropriate, by supporting mandatory application.
4. Support training initiatives in safety regulatory, and rating systems to promote policy making capacity, particularly in low- and middle-income countries (LMICs).
5. Promote the use of safer digital technologies by hospitals and healthcare providers in both the public and private sector. This will include competitions and awards for promoting best practice.
6. Recognise achievement in safety, innovation in safety-related technologies and practices, and products through a global awards scheme.
7. Target to halve preventable digital healthcare deaths and injuries.
8. Support improvements in evidence-based information to inform patients and staff about the performance and safety of digital healthcare products.

Car engineering is a professional field, but, as this booklet has shown repeatedly, digital healthcare isn't. Digital healthcare urgently needs to transform its culture. Digital healthcare needs to radically improve so frontline staff don't have to learn to cope with its limitations. Thus, as digital healthcare is way behind car safety, we have added some important points that NCAP didn't need to mention ...

9. Introduce and require safety measurement and labels. Like NCAP, make assessments public.
10. Provide digital healthcare qualification frameworks so digital competence can be measured and forced to improve. Healthcare providers, manufacturers, and regulators must be able to confirm staff are appropriately competent and accredited to work in posts responsible for development, procurement, investigations, or leadership, etc.
11. Once there are qualifications, it must be a requirement that developers are appropriately accredited. Note that drug development routinely uses post-doctoral scientists, so where are all the post-doctoral software engineers developing and testing digital healthcare systems?
12. Digital qualifications and professional accreditation must be developed and regulated through a new competent professional body with external oversight. It's an international problem, and needs international collaboration.

Fix IT page 467

We need a new organisation, preferably with statutory powers, to set standards and lead a Digital Health Improvement Program.

THINK LIST

What can we do?
Share this booklet.
The list below highlights key points to think about.
Nothing will be effective if people don't recognise the problems that need solving.

1. **Don't get sucked in by new technology. Be curious and think critically. (p. 3)**
2. **People who develop for digital healthcare need appropriate software engineering qualifications. This should be required by law. (p. 5)**
3. **You depend on digital. Are the systems you rely on dependable? Were their developers qualified to build safe systems? (p. 7)**
4. **Be slow to blame staff. Ask — could computers, devices, and records be unreliable? (p. 9)**
5. **The legal presumption that computers are reliable, and similar mistaken assumptions in internal investigations, need correcting. (p. 11)**
6. **Check the cybersecurity compliance of systems you rely on. Frontline staff should not have to learn how to compensate for poor system design. (p. 13)**
7. **Understanding how design error, Human Factors, and the psychology of error combine with poor digital design is essential for incident investigators. (p. 15)**
8. **Know your limits. If you don't know your limits, and aren't aware of them, you aren't safe developing digital systems — even simple Excel spreadsheets. (p. 16)**
9. **There are things you can do to prepare for any investigation. (p. 17)**
10. **Digital conceals complexity. Computational thinking is the right approach to manage complexity. (p. 19)**
11. **Digital healthcare should be evidence-driven. More digital healthcare research should be done, and the research needs proactively adopting. (p. 20)**
12. **We need safety ratings to inform procurement. Manufacturers will then make safer and safer systems as they try to out-compete each other once the quality of digital products is made obvious. (p. 21)**
13. **We need a new organisation, preferably with statutory powers, to set standards and lead a Digital Health Improvement Program. (p. 23)**
14. **Up to date qualifications are a key way to check what people do and don't know. (p. 26)**

FIX IT PRIZE

This booklet is based on Harold Thimbleby's book *Fix IT*, which won the British Medical Association's General Medicine book award. The judging panel said:

> *Fix IT* is such an important book. Our ability to help patients is so reliant on IT and digital solutions. It has the broadest appeal and has achieved something quite impressive. It is not just medically-focused in presenting solutions. A real strength is that it takes examples from outside of healthcare and translates them into healthcare. It should be read by all healthcare staff.

Harold Thimbleby, *Fix IT: See and Solve the Problems of Digital Healthcare*, Oxford University Press, 2021.

The Fix IT prize

In collaboration with the Royal College of Physicians, we have launched an annual prize for digital innovation, **The Fix IT in Healthcare Prize**. The prize rewards excellent innovations in digital healthcare, which help solve the problems described in this booklet, particularly digital initiatives that have improved, or promise to improve, patient safety and staff well-being within the NHS or healthcare internationally.

More details are available from the Royal College of Physicians web site, www.rcplondon.ac.uk, under Funding & Awards.

WE DON'T KNOW WHAT WE DON'T KNOW

In 2017 I built a garden office. I am good at DIY, but I made a rookie mistake that cost a lot. I had paid some builders to lay a concrete foundation, which I watched them do — I even helped them make sure it was the right size and perfectly level. When the concrete had set, I built the office, complete with its roof to keep the rain out. I then started to fit the details — and discovered the office was like the Leaning Tower of Pisa. The doors wouldn't fit. Everything was leaning. You'd think I'd have noticed? The concrete base was in the garden, surrounded by trees, and there was nothing upright or level to compare it to.

When I'd helped the builders lay the concrete, they made a frame the right size, we all checked it was level, then they filled it with concrete. I'd watched them do it!

Unfortunately, when they tamped down the concrete, the frame sank down in one corner. So the concrete base was finished off sloping down to that corner. With foundations on a tilt, the office I'd built wasn't square.

I had to dismantle it, and get the builders back to re-lay the concrete more carefully. Then I had to start again.

There's an important Human Factors lesson here. If you aren't trained, or you get distracted, you may miss a critical step in a procedure. I made a mistake *because I didn't even know I was making a mistake*.

Fortunately, there are ways to avoid or correct mistakes. A first step is to recognise our limits, for instance no self-aware DIYer would contemplate building a tower block.

Like DIY itself, DIY programming is very popular. It's very easy to start a DIY digital health project, but if we make a mistake we don't notice and so don't fix it, our code may harm thousands and put staff in prison.

? How can we recognise when we might hit our limits, and know when to say no?

The next column sets out the harsh solution: if you haven't got professional software qualifications, you have *no idea* what your digital health limits are — whether you are a regulator, developer, investigator, or clinician.

The tragic Down story (page 6) centred on a calculation that went wrong. There were at least two faulty assumptions the programmer hadn't realised they were making. First, they calculated the mother's age wrongly (the Millennium Bug), and, secondly, they used a negative number in a formula that was supposed to take a mother's age — which should be a number between, say, 10 and 80, and certainly not a negative number.

In programming, these types of assumption are called *preconditions*. A professional software engineer will prove all preconditions are met, but even if there was a proof, the code may have unnoticed typos or other problems (like hardware faults) that cause bugs, so their program will also have runtime assertions to check all the preconditions are met every time it runs.

For a software engineer, it is second nature to work like this. There are many more essential techniques, including Correct by Construction (CbC, developing programs directly from mathematics), code review (working in teams), Human-Computer Interaction (HCI), and rigorous testing that are central to safe system development. Someone in your coding team might ask, "How did you check the preconditions?" in exactly the same way somebody in a surgery team might ask, "Did you count everything at the beginning and end of the surgery?" to help avoid an inadvertent retention in a patient.

If you haven't been taught how to do a professional job, you have no idea what you do not know, and, worse, you are likely to be over-confident because you are unaware of your limitations. On the other hand, if you have been taught and examined, you will know what you scored in your exams, and hence how much — or how little — you know. You may want to learn some more!

To make matters more complicated, software is changing rapidly — blockchain, AI, cyber, robots, VR, implants, ... are all rapidly developing areas. Safe software is also improving very rapidly. So not only do you need good software engineering qualifications to work reliably in digital healthcare, you need qualifications that are constantly updated. *Fix IT* page 367

Up to date qualifications are a key way to check what people do and don't know.

FURTHER READING

- Clinical Human Factors Group (CHFG), https://www.chfg.org
- Healthcare Safety Investigation Branch (HSIB), https://www.hsib.org.uk
- Institute for Safe Medication Practices (ISMP), https://www.ismp.org

  These three organisations are wonderful resources.

- Marshall, P, Christie, J, Ladkin, PB, Littlewood, B, Mason, S, Newby, M, Rogers, J, Thimbleby, H, & Thomas, M. "Recommendations for the probity of computer evidence," *Digital Evidence and Electronic Signature Law Review*, volume **18**:18–26, 2021. DOI: https://doi.org/10.14296/deeslr.v18i0.5240

  If you're involved in a court case, your lawyers will find this free online article helpful.

- Mason, S & Seng, D (editors). *Electronic Evidence and Electronic Signatures*, OBserving Law, 2021. https://www.sas.ac.uk/publications/electronic-evidence-and-electronic-signatures

- Ferres, V (Chair), Caddy, C M, Evans, J, Falconer Smith, J & Jones, R. Northern General Hospital NHS Trust, *Report of the Inquiry Committee into the Computer Software Error in Downs Syndrome Screening*, Report submitted on behalf of the inquiry team to the Chief Executive of the Norther General Hospital NHS Trust and The Regional Director of Public Health. Undated.

- National Cyber Security Centre (NCSC), https://www.ncsc.gov.uk

- NHS Just Culture, https://www.england.nhs.uk/patient-safety/a-just-culture-guide
  See more details of Just Culture on page 17.

- PWC. *Conti cyber attack on the HSE* (redacted), Independent Post Incident Review, Commissioned by the HSE Board in conjunction with the CEO and Executive Management Team, 2021. https://www.pwc.ie

- Thimbleby, H & Cairns, P. "Reducing Number Entry Errors: Solving a Widespread, Serious Problem," *Journal Royal Society Interface*, **7**(51):1429–1439, 2010. DOI: https://doi.org/10.1098/rsif.2010.0112

- Thimbleby, H. *Fix IT: See and Solve the Problems of Digital Healthcare*, Oxford University Press, 2021.

  Page references to *Fix IT* are provided throughout this booklet. *Fix IT* itself has over 500 notes and references to further material.

- Thimbleby, H. "NHS Number open source software: Implications for digital health regulation and development," *ACM Transactions on Computing for Healthcare*, **3**(4):42:1–42:27, 2022. DOI: 10.1145/3538382

  NHS Numbers are ubiquitous and should be easy to program, so if you're skeptical about how bad NHS software can be, read this paper. It discusses the limitations of standards like DCB 0129 and DCB 0160.

- Wallis, N. *The Great Post Office Scandal*, Bath Publishing, 2021

ACKNOWLEDGEMENTS

See Change (an anonymous funder), Robin & Marianne Anker-Petersen, Nicholas Bohm, Bev Littlewood, Lydia Soar, Anne & Martyn Thomas, Sarah Toombs, Dave Williams, ...

Patient Safety
Stories for a digital world

Building on the award-winning book *Fix IT*, this quick-read booklet makes the stories of the promises and risks of digital healthcare engaging and accessible to all healthcare and related staff — nurses, doctors, investigators, regulators, digital developers, managers, and politicians.

This important booklet includes recommendations on what to do if you are blamed for a safety incident related to digital technologies.

Temporary picture

Prue Thimbleby is an artist and the Digital Story Lead for Swansea Bay Health Board in Wales.

For the last twelve years Prue has been recording people's stories about their experiences in healthcare, particularly stories of when things have gone wrong.

The stories are edited and co-produced with the storyteller, so that the storyteller has control of the story. The end result is a carefully-edited short (1 to 4 minute) video that powerfully tells the person's first-hand experience of healthcare. Stories have informed decision making, and improved services across the NHS.

The methodology Prue has pioneered is known as the Swansea Bay Method of digital storytelling. It has become policy for collecting feedback in the NHS in Wales, as well as being adopted widely across the NHS in England, as a result of Prue's training courses.

Prue is Council Member of the Arts Council of Wales.

More details are available on her web site, http://www.artsinhealth.wales

Prof Harold Thimbleby is See Change Fellow in Digital Health, based at Swansea University, Wales. Harold is a popular speaker, and has been invited to talk in over 30 countries.

Harold won the British Computer Society's Wilkes Medal. His most recent books, *Press On* and *Fix IT* have won several national and international awards.

Although a professor of Computer Science, he is an Honorary Fellow of the Royal College of Physicians, the Edinburgh Royal College of Physicians, and the Royal Society of Arts. He's also a Fellow of the Learned Society of Wales.

Harold has been a Royal Society Wolfson Research Merit Award holder and a Leverhulme Trust Senior Research Fellow, and he is 28th Gresham Professor of Geometry. Harold is Expert Advisor on IT to the Royal College of Physicians, a member of WHO's Patient Safety Network, and an advisor to the Clinical Human Factors Group, and the UK Medicines & Healthcare products Regulatory Agency, MHRA.

More details are available from Harold's website, http://www.harold.thimbleby.net