The cutting edge of HCI

Harold Thimbleby
harold@thimbleby.net

Harold Thimbleby wrote Press On (MIT Press), which won the American Association of Publishers award for best book in information sciences.

Meteorites were our first source of iron, thousands of years before the Iron Age when the strength, durability, sharpness, flexibility, and ductility of iron started to transform society on a large scale. Although knives and swords were immediate applications of the technology, it took another ten thousand years or so to refine the surgical scalpel, even though in principle everything was already in place to do so much earlier.

A scalpel is just an interactive lump of iron, but developing the full social context to make sense of it, to make the tools to make it, to make it work dependably, to make it a mature tool rather than a toy, to ensure users had the knowledge of germs and toxic contaminants, took the time; millennia, in fact. The build-up of social communication and shared knowledge so that enough people knew what scalpels were all about was part of the delay. Concepts like asepsis didn't have much use until surgery got going, and surgery didn't get going until it was successful on a big enough scale for enough people to see, learn and disseminate its basic concepts. Today scalpels are a small part of a large, intricate socio-technical system.

More complex uses of iron make the point even clearer. The car would be useless without a road network, and our pervasive road network pretty pointless without cars – in fact without lorries we couldn't have made the road network in the first place. Cars and lorries would be dangerous without a complex web of social conventions, which has taken at least a century to refine. Until roads meant that we could live far apart, we didn't want to travel so far; we certainly didn't need to commute.

Effective technology doesn't just happen. Scalpels are pretty good; now cars have a way to go; and computers, well ...

Emergency departments in hospitals use whiteboards to make notes about patients. In some hospitals these whiteboards have been computerised, but often with disappointing results. One of the iconic images of the disappointing failure of computerisation is a hospital room with ten scrubbed-up clinicians all standing around and using a conventional whiteboard, but with a computer screen opposite it, intended to replace it, being totally ignored.

It is not too far-fetched to see this new interactive computer technology being about as ineffective for today's clinicians as a scalpel would have been in Pompeii. A scalpel made then (around AD70) would have been a disappointing and misunderstood device: it would have been poorly designed, people would not know what it was, it would be dangerously fragile, and anyway, as used, it would have spread more disease than it cured.

It isn't such an extreme analogy.

To be effective, an ordinary whiteboard requires the supportive confluence of many factors: its users have to be able to read and write and understand the nature of permanent and temporary pens (and that some pens don't work at all on whiteboards); its designers have to understand that whiteboards must be smooth, robust, wipe-clean, sufficiently large, rigidly fixed on a wall. The physicality of the whiteboard has to match the task that it is being used for. There is a huge amount of tacit knowledge (such as the invention of colour codes to represent patient data) that is developed in concert with a deep understanding of the whiteboard/pen/writing/reading technology.

The point is that a whiteboard looks so easy to use, that surely a computer could only do the job better. Whiteboards are such primitive technology compared to computers! IT is, after all, so much more flexible and modern! With IT, you could put the whiteboard on the Internet and a clinician could view it from the other side of the world; you could add all sorts of other useful information about patients, from allergies to religious issues. You could work out how much treatment was costing, or you could track waiting times. Wow. E-wow.

We forget that to use a whiteboard requires skills that take the best of us maybe a decade to learn, and even then a few people never read and write reliably. I, for one, was the despair of my teachers, who thought that I would grow up and be a break in the continuity of civilisation.

The skills the successful whiteboard manufacturer needs are even rarer and harder to define. And to create the social context where the whiteboard happens to work so well took centuries, if not millennia, to sort out. When a whiteboard "just appears" in a modern hospital emergency department, we easily overlook this entire tacit socio-technical baggage.

Why does an expensive computer system fail so magnificently compared to a cheap sheet of plastic stuck on some mixture of epoxy and recycled woodchips?

Why does the hope of interactive systems continue to triumph over frequent frustration and failure? Why did anyone want to spend thousands of pounds on a small, unreadable display that nobody knew how to use, that would cost thousands more in wiring, that needs technicians and a maintenance contract and a backup system to keep it working – when something perfectly adequate was already working for a fraction of the cost? If the computer system breaks down, as it will, everything will grind to a halt, but if the whiteboard broke down (even the very concept boggles the mind) anybody would know how to work around it. They could write on the wall if somehow the whiteboard broke. Whiteboards don't even need rebooting, and if there is a power failure (and for whiteboards, it would only matter at night) a neat piece of technology discovered well before the Iron Age would get them working again.

Whiteboards look so simple that it must surely be simple to automate them. Indeed if you automated them properly you'd end up with something pretty much like a whiteboard, and there wouldn't be many advantages in that. If you don't automate them properly, you end up requiring a vast amount of unfamiliar tacit knowledge that nobody has.

Complex ways of failing are not the only problem with computerised whiteboards. They also can't be appropriated, extended, modified; it's not just that they don't fit in with deep social knowledge about use, they don't integrate with the many other technologies that do so well.

If not in hospitals, then, computers have been stunningly successful in some areas. Consider mobile phones, computer games, the Internet. Computers have also been stunningly successful in popular culture and in science fiction. They've even been successful in some mundane areas like payroll, although their success in finance seems to have been literally over-sold, as the reselling of complex financial instruments recently showed.

The point is: their stunning success in certain areas is no predictor of their success in other areas. In fact, it would be more truthful to call their "stunning success" anywhere an "accidental success". Mobile phones weren't planned to be so successful, and while text messaging became an unexpected success, many ideas failed terribly – but we can't recall what they were, as not many of us saw them in the first place! Fortunately, some things fail really quickly.

Not only are we excited and fooled by narrow success, it suits powerful interests to keep us excited and fooled. A whiteboard manufacturer doesn't have good profit margins and competes against plenty of other suppliers. In contrast, a "computerised whiteboard" supplier can sell an unfinished, unpolished bit of technology with huge margins and, moreover, lock the purchaser into a complex contract, to say nothing of paying for a training programme. Since anybody who can make computerised whiteboards can also make office information systems and lots of other stuff, they aren't going to fail quickly enough if they have one rubbish product. Badly designed whiteboards – badly designed interactive stuff – are going to be around for a long time.

You can look at a damaged conventional whiteboard and see at a glance if it won't work well; it's transparent, honest technology. But you can't assess a half-finished computer system and put a sensible price on it, predict how much its under-performance or errors will cost the people who try to use it, or even come up with reliable workarounds so you can stay working.

And this is where HCI comes in: to assess and understand how things work so that insights can go back to designers to improve the next generation of systems, and so that insights can go back to the rest of us who have to decide what to invest in to make our lives more effective, fun and worthwhile.

Some people in HCI have to cope with messes; there are indeed people studying hospital whiteboards, for example. Some hospitals need all the help they can get! But HCI must not confuse studying problems, fascinating as they may be, for the larger and more strategic responsibility of avoiding them in the first place. One hopes that HCI will do more than understand or improve specific situations (for that is usability, not HCI) and be able to generalise insights into a transforming science. As the examples above made clear, the real contribution of HCI isn't knowing details like when voice input is better than a pie menu, it's contributing to the whole socio-technical context: helping designers use better processes, helping technical authors be honest, helping procurement choose wisely, helping managers hire competent programmers, etc. In short, helping everybody match the task and technology synergistically. I hope, putting more effort into defining good technology than studying the consequences of bad technology – how a whiteboard fails is much less useful knowledge than how to make a better one. That it failed is one thing; that anybody thought it would succeed is more interesting; that nobody (or not enough people) who developed it had been on an HCI course is a disaster. Good HCI wasn't there for the people who needed it.

It would be tempting to digress into the nature of reliable knowledge that HCI should aspire to so that it is effective in this undertaking, but that is a well-worn discussion (the philosophy of science) that is not about HCI, computers, human factors or users specifically.

Instead, the thought I want to leave you with is this: with computers, clearly, we can and have changed the world; with HCI we should be aiming to change it for the better, and, let's hope, doing so a good deal faster than those delaying interests that thought the most profitable use of iron was the sword and not the ploughshare or scalpel. Don't think that understanding usability problems is going to be as radical as strengthening the science behind HCI so that it has wider, faster and more reliable application.