Itch

A short story by Trevor Hopkins

Do you get an itch you can't scratch? No, not *that* kind of itch!

You know how it is. You get an itching, tickling sensation, somewhere in the middle of your back, and you can't quite reach to just the right place. Perhaps it's between your shoulder-blades, or just below, or to one side. Or if you think you can reach, it's never entirely satisfactory wherever you scrape or rub. Do you want to know what causes that itch, the one you just *can't* seem to scratch?

I know the reason why you itch. If you're sure you want to know too, read on.

I'm working as a Research Assistant. You know, one of those underpaid and overworked kids with lank hair and poor complexions to be found in some numbers in their natural environment – the quieter and darker corners of the science faculty buildings. The faculty itself is part of one of those red-brick Universities which was instituted in an act of Victorian philanthropy, and which has grown over time almost organically. The Uni has gradually displaced the back-to-back terraces and narrow alleys that surrounded it with newer buildings which were probably supposed to be soaring white edifices of glass and stone, but seem to have ended up as irregular piles of water-stained grey concrete.

Like Mycroft, my life runs on rails. During the day, I try to find enough time to make a dent in the seemingly endless task of completing my PhD thesis, between bouts of sleeping and eating from the nearby takeaway kebab shop known affectionately as the 'Armpit'. I spend the minimum possible amount of time in my room in a rented house I share with several other postgrads – which is just as well, since it is cold, squalid and damp.

At night, I'm working on computer models of brain function – a task as large and complex as the Human Genome project, although we're a long way off that kind of successful completion. This is one of those crossover subject areas between AI and Robotics (which has been the Wave of the Future for more decades than I've been alive)

and Bio-informatics (sponsorship home of the big pharmaceutical and healthcare companies). Basically, some of us have finally realised that we really don't know enough about creating smart systems – we need to know more about existing intelligences before it makes sense to attempt to build artificial ones.

Of course, brain function mapping has all sorts of potential spin-offs, which is why Big Pharma and the healthcare consortia are interested in what we do. So much of human behaviour is determined by our hard-to-predict reactions to external stimuli, and there's so much we could do with a deterministic model of the machine between our ears – everything from improved anti-depressants (which is a pretty big market these days) or even a better contraceptive with no side-effects. Yes, ladies, you might just be able to *think* yourself not pregnant!

Selling these big ideas to the big companies, and gathering in the resulting big research grants, is of course the responsibility of my university supervisor and his professor, leaving me the menial task of actually making the technology work.

So, I'm steadily fumbling my way towards constructing a highly-abstracted model of total brain function. It has to be a hugely simplified abstraction – even the immense supercomputer in the basement (supplied at an extremely cut price by Big Blue, who really know how to woo the Big Pharma marketplace) was theoretically capable of representing only a tiny fraction of human mental activity.

Really, I'm refining a nearly automated process. I've been developing a suite of programs, including a library of rapidly-reconfigurable heuristics, which is capable of a statistical analysis of a huge number of brain scans. We've a library of recorded scans from NHS hospitals all over the country, all completely anonymous of course, as well as access to the results of stimulus-response experiments from all over the world. With static, structural information available in increasingly detailed form from CAT scans and the like, and dynamic information from the experiments, there's a wealth of data in there which just needs a structure to pull it out.

So, my heuristics take the raw brain function data, map it to a set of conceptual ideas of brain function, and then compile it into an abstracted executable model in a form that can be executed directly on the thousand-odd processors of the machine in the basement.

In short, I've built a brain capable of being run on a supercomputer. You can't really tell what its thinking, or even if it *is* thinking in any real way, but you can tell if the model's responses to stimuli correspond to the measured responses in a real brain. There's just enough complexity in the model to show genuinely emergent behaviour and detectable emotional reactions.

Of course, this takes vast amounts of computer power, both to compile the model itself and to execute it. It takes an hour or so of all those processors crunching away to simulate the effect of five seconds worth of what I can loosely call thinking.

Naturally enough, most of this work is done in the middle of the night, when no-one else wants to use the machine. A few uninterrupted sessions in the wee hours are exceptionally productive, when the building is dark and quiet. The whole process is directed from the networked workstation in the corner of the office I share (if I was ever here during the day) with two other RAs and an indeterminate number (it seems different every week) of research students.

The whole process is trial-and-error, of course, which is why it takes so much time. The evaluation of each compiled model (a couple of hours on the supercomputer) is based on stimulating its simulated inputs. (Say that without spitting!)

Now, a large part of our brains is associated with processing optical inputs – there are other inputs as well, of course, but we are, fundamentally, visual creatures. So, part of the model itself, one of those conceptual ideas of brain function I mentioned, involves stimulating the optic nerves and modelling the corresponding movements of the eyes themselves. This coordination of eye movement and the inputs from the smallish number of high-resolution optical sensors in the retina is one of the novel features of this model, and it seems to successfully overcome some of the limitations in previous attempts to build a truly effective visual parser.

It's well-known that we use only a small fraction of our brain. Actually, that's not really true, more an urban myth. More sophisticated measurements and less intrusive techniques has allowed recent experiments to detect neuron dynamics in regions of the brain previously thought to be redundant. Still, there do seem to be some areas with no discernable purpose, and part of the research is to find out more about unused brain cells.

Basically, I showed pictures to the model. Some of these came from a library specifically for this purpose, but I found I got some interesting reactions, and in particular some dynamic behaviour in regions thought to be inert, by using images with distinctly emotive contexts. Some images were already available online whilst others I simply scanned using the multi-function printer-copier down the hall.

All was going well until I started showing the model pictures of naked people. Look, OK, this is the kind of thing you do when you're working all alone in the middle of the night, at a task which requires occasional flashes of insight, a few minutes of concerted effort and several hours of boredom. Besides, I knew about this collection of well-thumbed magazines hidden away in the back of the filing cabinet.

Of course, I expected some emotional reactions – perhaps some analogue of prudery and embarrassment in the higher regions, and some pretty direct sexual responses in more primitive areas. What I actually got was a curious mixture of disgust and loathing, even fear, and a distinctive activation of the 'fight-or-flight' reaction. If it was a real person, it would be feeling some horrific combination of stomach-turning revulsion and stomach-knotting fright.

I just had to investigate, although I've now come to seriously regret that decision. It's fairly easy to find out what part of an image the model is concentrating on, since it is, in essence, moving its eyes as it scans and comprehends the scene in front of it. I'm sure you can guess the body parts I had expected to attract. I was wrong. Over the course of an hour's run, the model's simulated eye movement ignored the external genitalia and various wobbly bits, and focussed almost entirely on a small area between the shoulder-blades.

You know, this might have been the moment I first started itching in that exact place?

I carefully checked for image defects and scanner problems, and found nothing. The model's reaction to images of people with their clothes on was unsurprising, and completely consistent with its response to other, less emotive, contexts. On closer investigation – yes, I really did download all those pictures from the Internet for scientific reasons – I found that the model would display plausibly randy reactions to pictures where the back and shoulders were not visible, but fear-and-loathing when presented with shoulder-blades.

One projected use of highly detailed brain models is truly effective hypnosis – the ability to remove compulsions and inhibitions, or even

be able to introduce them artificially. You can imagine the government and military wetting themselves thinking up ways of using that capability.

So, my initial hypothesis was that the model had somehow gained an artificial neurosis, produced as some obscure reaction to an anodyne part of the human body. These kinds of discrepancies between modelled and real-life behaviour are always interesting, and often a fruitful source of material for papers to be published in some of the more obscure journals. Oh, and of course it adds to my professor's credibility in the never-ending pursuit of sponsorship money.

My objectives were two-fold: first, to reduce the variables, to avoid any side-effects of image coding techniques or copyright-tracking steganography. For this purpose, I captured an image of me, from the back, and wearing no clothes. I borrowed a high-resolution digital camera from the image-processing labs on the next floor down, and used the most loss-less image encoding format I could identify. The single picture took up a substantial fraction of my personal disk space quota. I even printed out a copy and blu-tacked it to the wall above my workstation.

My second objective was to present the stimuli and the model's reaction in a way that was comprehensible to mere humans. I set about writing a new program to extract an image of how the model itself perceived the scene it was viewing. This took a lot of programming, and I sat up over my workstation for several nights until the new interfaces began to show signs of working.

During these few days, I found myself neglecting my write-up and sleeping even less than usual; inevitably I was compensating by eating even more of the blisteringly hot kebab-and-pitta-bread concoctions from the 'Armpit', washed down by alarming quantities of caffeinated cola drinks.

Finally I was ready for a full test run. Sitting at the workstation, I reloaded the most recent model, and hooked up the new visualisation software, then typing the few commands which started the model's reaction to the image of my back.

I'd displayed the evolving picture of artificial perception in a window I'd placed in one corner of the screen. It showed a desperately low resolution at first, with each pixels worth of enhancement being painfully computed as the kilo-engine processor in the basement ground away.

Eventually some kind of comprehensible picture began to emerge from the twin mists of simulated perception and digitised noise. Frankly, I was utterly horrified. The details that emerged showed some kind of growth, a green bump embedded in my own skin between my own shoulder blades.

Somehow, my own inherent perception changed at that moment. I've heard that expression about 'scales lifting from my eyes', and that was exactly what happened. My picture, stuck to the painted breezeblocks above the workstation, seemed to shimmer and twist, a green blob appearing before my eyes and between my shoulder-blades.

Opening one of the filing-cabinet magazines showed me pictures of bronzed muscular men and compliant young women, all with one – or sometimes more than one – of the green appurtenances protruding from their backs, just where the model showed they would be.

I rushed to the gents bogs and lifted my shirt, looking at my own reflection in the rather grubby mirror over the cracked washbasin. There it was, a bright virulent green, like a really ripe green pepper – a bell pepper or capsicum – somehow seamlessly merged with normal pink skin on my own back. In the mirror, I could see a slight sense of movement, somehow pulsating gently like a TV special effect from an early edition of *Doctor Who*, its movements distinctly out of sync with my own breathing and heart rate. I was heartily sick, there and then.

I think they're some kind of symbiote, or more likely parasite. They grow on people, on everybody, their roots digging deep into our bodies. My best guess is that the growths form links into the spinal column and produces some kind of hypnotic effect in our brains which prevents them from being seen. Somehow, we all share a worldwide neurosis, an induced inability to see what quite literally sits on our own shoulder.

I've been looking at these things for several days and nights now, not sleeping much. Now, I can see them everywhere, even detect their presence under tee-shirts and fleeces. Everyone has at least one and some people – particular very slender and attractive people – have several. Perhaps the physical drain of keeping two or three of the parasites alive from your own bodily resources means you have no excess fat – and the added induced neurosis that exceptionally well-inhabited hosts are both thin and beautiful.

What I still can't work out is why the growths are not hidden entirely within our bodies. My best guess is that they are some kind of

plant, and they can't quite get all the nutrients they need from us directly. So, they must retain some kind of vestigial photosynthesis, to produce some vital trace compound not available from our own blood streams.

Just for my own reference, just a way of hooking them to a name, I've taken to calling them Monkey Plants, after the expression 'a monkey on your back' in that modern sense of a serious problem that just will not go away.

I've spent some time thinking about how to remove them, and what would happen if I did. My back seems to be itching all the time now, and I know what's causing it, and I'd dearly love to rip the offending growth from my skin.

In fact, I'm not even sure that they even can be removed. I can see the Monkey Plant on my own back, but I can't touch it – not my own, not other peoples. Even though I can see exactly where it is, I can't control my own hands, or the movements of my body, to actually press my fingertips against its surface. There's something deeper in the hypnosis, something at a detailed level that my computer model won't let me reach, which prevents the physical contact. Another one of those supposedly inert regions of brain cells kicking in, I expect.

I'm not a parent, and may very well never be, but we've all heard stories of babies crying incessantly, inconsolable despite the best efforts of their increasingly fraught mothers. It must be incredibly painful, the initial infection before the first of these things has fully integrated itself with the spinal cord. A baby can't move in a coordinated way, or communicate; it has no way other than bawling to show the agony it is enduring.

I think the infections move from person to person, with some kind of seeds or spores being transmitted from parent to offspring, and growing and living with us for all our lives. Imagine our bodies aging, wearing out, drained by the incessant physical demands of feeding the things. I've seen old people, hunched and feeble, bent nearly in two by the Monkey on their backs.

These things don't think, in any way we understand the term. But they have desires, needs – they want to grow more, and the more people there are, the more they can grow. I suspect this endemic infection has pushed us, our society, in certain directions – to live in large groups, in groups and cities, and to alter our environment, our world, driving our evolution, making us invent technologies to give us

the resources to support more bodies – just so that they can reproduce more.

I think the reason human beings are taking over the planet is because the Monkey Plants have taken us over.

I believe there's yet another mechanism that the plants have evolved over the millennia. If you try and talk about them, the growths on your back, you are comprehensively ignored. Not disbelieved, just ignored, as if you had said nothing at all. I've showed my results, the pictures, to my supervisor, and he just changed the subject back to the next round of grant submissions – no real difference there, then. I've tried to engage some of the other postgrads in conversation, even buying the pints in the back-street pub we occasionally visit; again, they just don't seem to hear what I say.

So, this is my attempt to communicate – to tell the world about this disease, this parasite, which is warping our bodies, and our minds, and our societies.

And, you know what, I just bet you won't believe me. Oh, you'll read my words, even declare that you completely understand what I've written. But that monkey on your back just won't let you believe, really *really* believe. It's just a story to you, isn't it?

3029 words

8 pages

26/04/2007 21:04