

## **Big Mother project history, 1983-2019**

### **Ancient history**

- 1983 -

I joined Sinclair Research Ltd (SRL), as a software engineer, on 4 July 1983, armed with a humble BSc in Computer Science from Queen Mary College, University of London. (I rented a tiny bedsit in Devonshire Road that was so cold my quilt would be frozen in the morning!) During my time at SRL, I worked on the QL, Spectrum 128, and Tyche computers, i.e. SRL's "mainstream" products.

Within just a few months of joining, I found myself being asked to accompany SRL chairman Sir Clive Sinclair and Martin Brennan (a more senior engineer in the computer lab) on various "AI day trips" around the country (great fun!) To this day I have no idea why I was chosen, but to all intents and purposes I had effectively become one half of Clive's then-embryonic in-house "AI team".

- 1984 -

Sometime in 1984 (if memory serves), Martin and I were the first two people from the Willis Road lab to move into Clive's new "MetaLab" research centre, at Milton Hall. Later, an entire Wafer-Scale Integration team was brought in from Burroughs in order to develop massively parallel CPUs, for AI purposes, but Martin and I were still the only two software guys on the AI team.

And so I found myself, in an office at Milton Hall, tasked with the job of developing (i.e. writing software for) an AI machine, to run (in the longer term) on WSI-based CPUs. I had literally a blank sheet of paper, and absolutely no idea what I was doing - that was my starting point.

However, the mindset within the SRL computer lab was of course to develop (blue-sky) products, and so my working assumption was that whatever I was developing (the 30 second brief I'd been given by Martin was not very specific on this point!) it would be some kind of commercial product.

And AI, in those days, meant full monty Artificial General Intelligence (AGI), as it is known today.

Apart from the purely technical challenge (which was a dream come true for any techie geek), the opportunities for massive societal benefit were immediately apparent. The idea that the project also had the long-term potential to create a near-utopian dream for all mankind, which was not really something that any mainstream computer project could boast, was massively appealing. That dimension made the AI project special in a way that no other project could ever match.

Anyway, I didn't know it at the time, of course, but this was the genesis of (what would eventually become) the Man-Made Minions (MMM) project:

- commercially motivated
- Artificial General Intelligence (AGI)
- near-infinite potential for longer-term societal benefit

Not really knowing anything about AI, I simply treated it as I would any engineering project: start with a top-level specification (goal) and recursively decompose it into progressively smaller (and more detailed) components, until finally every component had been concretely implemented.

It took me several months (of contemplating my navel, and reading everything I could about AI) to decompose "AI" into "reasoning, planning, and learning", all connected via a "knowledge base".

- 1985 -

It was immediately obvious that this was way too much to develop all in one go, and so I decided to pick one of "reasoning, planning, and learning" to do first, as an initial product, and then, longer term, to gradually extend "one cognitive dimension at a time" until full AGI had been achieved.

I picked “reasoning” as the basis for the first AI-based SRL product, not because I knew anything about it (I didn’t), but because (deductive) reasoning is basically mathematics, which had 5,000 years of history behind it, and therefore something to start from, whereas (in 1985) “learning” and “planning” were still dark matter. Problems with the source code for the new Spectrum 128 editor inspired me to select “formal program verification” as the first, reasoning-based, AI product. This was also very promising commercially because software bugs had been the greatest problem in computer science for 40 years, and so mitigating it would be a massive contribution to the field.

Naturally, I knew nothing about formal program verification either...

But then the SRL AI software team acquired its third member - Aubrey de Grey, who *did* know about formal program verification. After some discussions, Aubs very memorably snapped the book he was reading shut and proclaimed “I know how to do it!” He tried to explain it to me and I tried to pretend to understand (I didn’t), and off he went programming his prototype verifier.

- 1986 -

I meanwhile was pulled off to work on the new Tyche computer (QL successor), and a few months later Amstrad bought SRL’s trademarks and current products (but not work-in-progress, such as Aubs’ verifier); a few days later Aubs and I decided, with Clive’s blessing, to spin off MMM.

And so, in 1986, Man-Made Minions (purveyors of artificial workers) was born; our stated goal being “profit through the humanitarian application of machine intelligence”. Clive somehow finagled a deal whereby Aubs was able to stay on at Milton Hall and keep using the VAX 11-780 computer on which he was writing his verifier, we (eventually) negotiated a licence deal with SRL in respect of the IP, and I got a contract job at Verran Micro-Maintenance in Surrey to pay for it all.

The overall MMM plan, of course, was basically as it had been at SRL:

- step 1: build narrow AI product (formal program verification)
- step 2: use narrow AI revenues to generalise narrow AI into AGI
- step 3: near-infinite potential for longer-term societal benefit
  - although we had no idea how the intended transition from medium-term commercial goals to longer-term philanthropic goals would work out (we just assumed that it somehow would)

And so, it would seem, we were set. Only it wasn’t quite as easy as we’d initially thought. Running even a 2-man company costs a lot more money than you might expect, and writing a formal program verifier takes a lot more time than you might expect. So it was quite a struggle. But we believed very strongly in the project, and in the long-term plan, and so we persevered.

In the end, it took seven increasingly-financially-difficult years just to produce a prototype.

- 1988 -

In August 1988, Aub’s verifier found its first *real* bug, i.e. one that had not been planted by Aubrey.

Also around that time, our friend Nigel Halse suggested “if you ever build your [AGI] machine, you should call it Big Mother”... :-)

- 1989 -

By 1989, Aubs could no longer continue to work out of Milton Hall (I think they finally worked out that they had squatters sucking vast numbers of CPU cycles out of their VAX, and we were asked to leave!) And so we had to buy our own IBM PS/2s running OS/2, and rent somewhere for Aubs to work, and Aubs had to port his verifier from VAX Pascal to Prospero Pascal. More expense!

- 1991 -

By September 1991, things were starting to look sufficiently promising that I started looking for external funding (early-adopter corporate sponsorship; specifically *not* VC) for the next stage.

## - 1992 - First financial implosion

1992 was a momentous year for us. Aubs finally had a working prototype that, among other things, could formally prove “bubble sort” fully automatically (finding all the loop variants and loop invariants, and proving all the theorems, fully automatically). At the time, no-one in the world could do this, and so, despite the fact that Aubs’ verifier was still just a prototype, and highly unstable, we nevertheless perceived this to be the technical milestone (that would allow us to “cross the chasm” and reach an SRL-style mass market) for which we had been so painstakingly striving.

As far as I know, in 2019, no formal program verifier in the world can formally prove “bubble sort” fully automatically (finding all the loop variants and loop invariants, and proving all the theorems, fully automatically). And so, incredible as it might seem, it would seem to be the case that Aubs’ verifier had a ~25 year lead on the 1992 state of the art. Of course, we had no idea at the time.

We held an MMM fundraising seminar at Peterhouse, Cambridge’s oldest college, pitching to about a dozen safety-critical companies such as Siemens Plessey Defence Systems (SPDS), British Aerospace, Rolls Royce, and others. However, due to the very “opaque” nature of the prototype verifier at that time, making it difficult to evaluate by an outside observer, we were told that what we were doing was all great etc but that “we didn’t actually *have* anything yet” (words that would haunt me for decades!) And so, with one exception (SPDS) we failed to find any company, even in the safety-critical world, that was willing to back the commercialisation of Aubs’ verifier (which needed re-writing as a commercial product). Eventually, even SPDS fell through.

That same year, the Verran contract came to a natural close, and so MMM’s finances completely collapsed. I remember having just £1/day to spend on food (FYI you can get ravioli on toast and a can of coke for that). It was the first time in my life that I was literally hungry, with no way of simply buying e.g. a sandwich to make it stop (I used to hang around the kitchen after my housemates had finished cooking and eat e.g. the potato crust on the inner edge of their saucepan - heaven!) Ultimately I was forced to move back to my mum’s in Northampton, which I naturally hated.

This was a very difficult period, and it would have been extremely easy to give up at that time.

But either you believe in your project, or you don’t. And AI, remember, is special - it’s addictive. Once you’ve had your sights set on creating a global utopia via AI, nothing can ever take its place. And so Aubs and I decided to carry on; to generalise the ideas behind the verifier in line with the longer-term AGI plan; to re-implement the prototype verifier (as a commercial product) on top of that generalised architecture; and to do it all by ourselves, without external help, if need be.

By this time, I had started to catch up a little with Aubs re the technical aspects of the project, and the technical (AI generalisation and formal-language-related) work now before us was something I was more suited to than Aubs. And so we swapped. I took over the technical side of things, and Aubs went out and got a job to support me while I did so. (Or, at least, to partially support me.) Desperate to break out of Northampton and back to Surrey, I registered with ~30 IT recruitment consultants simultaneously (in order to supplement Aubs’ financial contributions), but due to my now bizarre work history I found it almost impossible to get a bite via normal channels.

And so here began (for me) two years of clinical depression due to extreme disappointment, frustration, financial anxiety, etc in respect of MMM and our difficulty raising funding, etc. For the next three years, I spent two nights a week sleeping in my car (in an M3 service station car park). The pain of clinical depression is such (only those who have experienced it will understand) that I also formulated detailed plans for suicide (cleverly incorporating a 12-hour cooling off period), in the event that I was unable to make it stop. As it turned out, those plans were never initiated.

- 1995 -

As part of the process of putting together a joint application for EU funding, I attended the *High Integrity Programmable Electronic Systems* workshop at Schloss Dagstuhl in Germany (not being able to afford a plane ticket, I drove from Northampton to Bavaria - the other attendees couldn’t believe it!) When our application failed, my co-applicant, Dr Wolfgang Halang (Wolfie) withdrew,

complaining that it was a lottery. Not wanting to give up quite so easily, I drove (again) to Brussels to meet with the relevant EU project officer, who was clearly impressed by my perseverance.

My depression now behind me, I was awarded a bursary to do the MSc in Computation at Oxford University (full disclosure: I wasn't actually awarded my degree), which as much as anything else was fantastic post-depression therapy. I took the opportunity to progress the MMM/AI project as much as I could while I was there (not just in respect of my MSc thesis topic, but I was also able to sit in on lectures, e.g. on denotational semantics, that weren't actually part of my course).

We also had some encouraging fundraising success, in the form of a small EU grant (second time lucky!) - this was basically a small grant to fund the process of applying for a larger EU grant...

- 1996 -

Post-Oxford, I was finally able to move back down to Surrey, but MMM's finances were still extremely fragile. This was the start of my 5-year "Petrol Station Food" period (during which I lacked the funds to shop in supermarkets etc, and so I had to do all my food shopping in petrol stations, paid for by credit card, but with no actual credit on it; in those pre-internet days they didn't check for amounts under £20, so I kept an emergency £20 note in my wallet just in case!)

- 1997 -

The EU project officer who I had driven to Brussels to see subsequently hired me as a technical expert (in advanced formal verification) for the European Commission, and thus, over the next 5 years or so, I was a formal reviewer for a number of Esprit projects. It was very little money, only a few days' work every six months or so, but it boosted my credibility a little, and also my morale.

I spent several months preparing an application for a much larger Esprit grant (1.6 million EUR split across a pan-EU consortium), but this was unsuccessful (Wolfie was right - it is a lottery!)

- 1999 -

By 1999, Aubs was working in the Cambridge University genetics department, and had become increasingly interested in biological longevity, so much so that his book *The Mitochondrial Free Radical Theory of Aging* led to him being awarded a Cambridge PhD in biology the following year.

- 2001 -

After almost ten years of searching for (non-grant, non-FFF) funding, we finally signed a deal (memorably, on 11 September 2001) with a small group of angel investors. It was only a tiny amount of money, drip-fed each month, and so the Sword of Damocles still swung precariously above our heads, but at least my Petrol Station Food days were finally over - hello Waitrose!

Given this relative stability, I was able to start actually coding the lowest levels of the emerging generalised AI architecture. Anxious to avoid the quality problems that had ultimately plagued Aubs' prototype, I developed my own bespoke software quality management tools, and heavily tested each module before allowing myself to move on to the next (some of the test sets took several weeks to run, and when they found a bug I'd have to run the whole thing over again!) I probably wrote around 30,000 lines of C during this period, but it was all very low-level stuff.

- 2002 -

In February 2002, having realised that raising money was likely to be a continual, never-ending process, I attended a week-long business planning course in Hamburg. One of the marketing experts, having grilled me for 2 hours about MMM's business plan, said that I was either a genius or crazy, but he couldn't tell which. I asked him if he could see the guy standing in the corner ;-)

- 2003 -

In 2003, Aubs' increasing interest in life extension led him to co-found the Methuselah Foundation, a US-based non-profit dedicated to extending the healthy human lifespan.

- 2005 - **Second financial implosion**

In early 2005 our angel investors could no longer continue with the monthly drip feed and, with zero safety net, cashflow buffer, or any financial headroom, MMM's finances immediately hit the wall for the second time, and I was again forced to move back to my mum's in Northampton.

- 2006 -

On this occasion, despite the obvious similarity of the situation, clinical depression did *not* ensue - I guess what doesn't kill you really *does* make you stronger! I tried absolutely everything I could think of to break out of the rut (and out of Northampton), but to no avail. By this point, my work history was now even more bizarre than it had been post-1992, and I was even more of an outlier in my own industry (software) than I had been before. And so I again found it impossible to find any consistent, reasonably well-paying computer programming work (which would of course have solved all our problems in an instant). Instead, all I could get was a minimum-wage temp job in a warehouse, packing boxes and doing other menial chores (except that when they discovered I could program they also asked me to build a spreadsheet - for no extra money of course). It was absolutely soul-destroying, and I only managed to endure it for about six weeks (never again!)

On a positive note, my 2002 business planning course finally paid off when we managed to secure a decent chunk of funding (spoon-fed rather than drip-fed) from a new investor, and, in the months that followed, further, smaller, amounts from a number of further investors (mostly friends).

- 2007 -

Once again, the Sword of Damocles had been temporarily returned to its sheath, but even so I did not yet feel sufficiently financially secure to move back to Surrey, and so I stayed in Northampton, even though I hated being there. My caution was ultimately borne out when, as funds were used up, replacement funds became harder and harder to secure. In the end, we were surviving on a few £100 here and there, and from week to week. It was simply not sustainable in the long term.

- 2008 -

Although it might not seem obvious from the above, interleaved with all the false starts, setbacks, dead ends, and personal and financial crises, useful technical work was actually being achieved.

In 2008, as Aubs had done in 1985, I finally had my own "I know how to do it!" moment. I could see, in my mind's eye, just as Aubrey had, a path to a technical solution to the problem on which I had been working (since taking over from Aubs in 1992). Not the solution itself, but a path to it.

What I had finally worked out was how to build the generalised AI architecture on top of which a commercial version of Aubs' verifier, and much more, could be built. In effect, I had designed the "inner cognition" part of an AGI, the "universe of discourse" of which was first-order set theory, and whose underlying cognitive processes were deductive reasoning and generalised abduction.

In principle, "full" AGI could then be achieved by building an "outer cognition" on top of the "inner cognition", effectively extending the universe of discourse out from first-order set theory to the entire physical universe. In 2008, I hadn't yet thought that far, but that was the longer-term step.

What this meant was that the architecture that I had devised would be able to support not just a commercial version of Aubs' verifier but also further products such as program synthesis (i.e. the automatic generation of computer software from a formal specification). The prospect of the latter was becoming increasingly important because, in the decade since the mass popularisation of the internet, the nature of the software industry had changed such that formal program verification

was becoming harder to sell to millennial coders than it had been to old school programmers. However, the potential market for program (similarly hardware) synthesis was virtually unlimited. So, technically, 2008 was a great year, but, financially, it was still a Sword-of-Damocles disaster.

- 2009 -

Despite my technical success (and I really did feel like a milestone had been passed), we still had the usual severe cashflow problems, and they weren't about to get any better any time soon.

And so, in January 2009, with our intermittent external funding having deteriorated to the point at which it simply wasn't sustainable, we decided to put the MMM project into cryopreservation in favour of a new, much simpler project with the potential to generate cash for the MMM project, at which point, of course, MMM could be revived, and the overall long-term AI/AGI plan restarted.

Given how limited our funds were, some kind of internet project was the only viable possibility, and so I devised, and started working on, an e-commerce project called MummyBear.

Also in 2009, Aubs co-founded the non-profit SENS Research Foundation in California. Briefly, Aubs' longevity research had identified seven cellular and molecular processes which together account for all known symptoms of human ageing. The effects of these seven processes accumulate throughout a person's adult life, ultimately becoming pathological and causing a breakdown of function, vulnerability to illness, etc. The basic SENS idea is to address each of these seven causes individually, and to fund scientific research into therapies designed to remove, repair, replace, or render harmless the cumulative damage caused by each one, effectively rejuvenating the human body, and thereby extending healthy lifespan, potentially indefinitely.

Thus life extension had become Aubs' (main) thing, and AGI remained (as always) my (main) thing.

Six months after starting work on MummyBear, so basically just as I was getting into my learning curve, my mum was diagnosed with dementia, and I became her full-time, 24/7/365 carer for the next six years, basically without a break. It was an unimaginably ghastly time, changing my mum's incontinence pants, feeding her liquidised food through a syringe, and once even finding her in a dehydration-induced coma, all amplified by illness-induced poverty. Naturally, the effect on the MummyBear project was catastrophic - I was lucky if I could get an hour's work done each day. But I persisted as best I could. Life had simply gone from hard slog to very very very hard slog.

- 2010 -

In 2010, following a day trip to a safety-critical conference in Bristol, I was tempted back onto the MMM project (basically, it's my true love, and thus extremely hard to give up, even temporarily).

And so I spent a year basically proving theorems of first-order logic by hand, looking for patterns that I could automate in the next part of the AI architecture to be built - the MMM theorem-prover.

It was a moment of weakness, but unsustainable, because nothing had changed financially.

- 2011 -

From 2011 I was back on MummyBear, and resolved not to return to the AI stuff again until the e-commerce project was operational, and earning revenues for MMM, which was the whole point.

And so, relative to 1986, the overall MMM plan had now evolved into:

- step 1: complete MummyBear e-commerce project
- step 2: use MummyBear revenues to revive MMM project:
  - step 2.1: complete generalised "inner cognition" AI architecture
  - step 2.2: build commercial narrow AI products on top of "inner cognition"
    - formal program verification
    - program synthesis
    - hardware synthesis

- ...
- step 2.3: use narrow AI revenues to build “outer cognition” on top of “inner cognition”
- step 2.4: near-infinite potential for longer-term societal benefit
  - although we had no idea how the intended transition from medium-term commercial goals to longer-term philanthropic goals would work out (we just assumed that it somehow would)

In other words, the AI/AGI project had now been generalised, and pushed down on the “stack”.

- 2012 -

Aubrey’s own mum had passed away in 2011, and in 2012 he was able to inject some cash into MummyBear from his inheritance. It was a godsend (and further godsend would later follow).

- 2014 -

After falling into a dehydration-induced coma, my mum was revived, and rehydrated, and moved into a nursing home, having been classified as end-of-life by her consultant. At the time, we all believed she had only a few days to live, but once installed in the nursing home she rallied.

Now having a little bit more time each day, the rate of MummyBear progress increased a little.

- 2015 -

My mum finally passed away in June 2015, by which point, having looked after her virtually 24/7/365 for six years, I was completely and utterly exhausted in every way imaginable. It took me quite a long time to work out how to proceed (all I really wanted to do at that point was sleep!)

- 2016 -

I needed to complete the MummyBear e-commerce project, per the overall MMM plan, but I also needed to fund it somehow. By now, my work history over the previous ~25 years, compounded of course by the six years that I had spent looking after my mum, was so off-the-charts bizarre that I was no longer a mere outlier in the software industry - it was more like I was off-planet.

I now estimated the probability of my ever being able to find programming work in a company that I didn’t own myself as zero. The alternative - minimum-wage work of some kind as per 2006 was unthinkable - and besides it simply wouldn’t have generated enough money, and/or allowed sufficient spare time, for the e-commerce project to be progressed at a realistic rate anyway.

And so the only realistic solution was to sell my mum’s modest house, which I had just inherited, and live off the proceeds, in rented accommodation, until MummyBear was up and running. After various costs, I was left with around £90k, which I calculated would fund the project for two years. I figured there was about a year’s worth of work still to do, which left me with a year’s headroom.

I had by this point been longing to move back to Surrey for some 11 years, but nevertheless rents in Northamptonshire are a lot cheaper, and I figured another year or two in the wilderness was no big deal. So I rented a lovely 200-year-old cottage, which I dubbed Rancho Relaxo, in the villages.

After a bit of a slow start, I buckled down to MummyBear, progress now greatly accelerated.

- 2017 -

Lots of progress was subsequently made, especially on the database and server side of things. But, of course, things were taking longer than expected, and after about 18 months I realised that I needed to raise further funds before my own ran out. (In reality - and it had taken me a lifetime to work this out - high-tech projects need more than just 1 or 2 people and a shoestring budget!)

And so I went through the motions of writing a business plan and pitching to angel investors etc, but, this not being my first rodeo, I knew in my heart that it would be virtually impossible to raise funds for another partially-finished software product (“you don’t actually *have* anything yet”, etc).

### - 2018 - **Third financial implosion**

Perhaps if I'd been a better salesman, in 1992 or in 2017, or if I hadn't made whatever strategic errors I'm bound to have made over so many years, or maybe if my mum hadn't fallen ill in 2009, then I'm pretty sure that I would have been able to develop MummyBear and/or Aubs' verifier to completion, but, in the event, that's not what happened. What happened was that I completely ran out of money in March 2018 (exactly as my spreadsheet had predicted would be the case).

On previous such occasions it had always been possible, with a little thought, to find a way to dig myself out. But not this time. Having put all my resources into the MMM project my whole adult life, I was now almost 57, with no property, savings, or other assets, no pension, and (because of my bizarre history) no realistic prospect of gainful employment in the normal software industry.

This meant that, once my Rancho Relaxo tenancy expired at the end of March 2018, with no money for rent, or movers, or storage, or any realistic means of earning a liveable income, and zero safety net, I would have literally nowhere to go, nowhere to put my stuff, not even my cat.

I couldn't see any way to avoid homelessness at this point. Not that it was anyone else's fault but mine, of course - I had understood the risks when I was taking them. But that was irrelevant now.

Years ago, I spent a sleepless night on the streets of London, and it was absolutely miserable. I really didn't want to end up sleeping under a bridge, or whatever, but I couldn't see any way of avoiding it, if not imminently then a few years down the line as I got nearer to retirement age.

And so I'm sorry guys but I basically gave up at this point. I don't believe that anyone could have tried any harder, or persevered more, than I had, but it was just not going to be possible for me, as per the MMM plan, to repay everyone for their contributions and support; it was not going to be possible for Aubs and I to derive any income, let alone personal financial security, from the MMM project; and at this point it wasn't even going to be possible for me to avoid homelessness.

When my mum was declared end-of-life in 2014, we had to decide, considering all the facts, what course of action was most in her best interest. It boiled down to a choice between a *longer* life followed by (a greatly increased probability of) a *bad* death, versus a *shorter* life followed by (a greatly increased probability of) a *good* death. We (her consultant and I) chose the latter. In January 2018, I was basically confronted with the same dilemma - and I made the same choice.

And so I decided to spend my last few months chilling with my cat in the lovely Rancho Relaxo, binge-watching Netflix and eating lots of Chinese takeaways, as well as taking the opportunity to socialise as much as I could, before taking my own life in a dignified manner of my own choosing.

Having developed my ideas for an "inner AGI cognition" entirely independently of the mainstream AI world, I had absolutely no idea if they contained anything worthy of merit. For all I knew, the "real" AI world would immediately dismiss them as laughably-amateurish and over-simplistic.

On the other hand, on the off-chance that there was anything useful in there then it would all be lost on the event of my death, and all my (rather extreme) hard work and sacrifice for 35 years would have been for nothing. And so I decided to write it all down, as best I could, in "How to build an intelligent machine", which I quickly threw together and published on Amazon (and I also took Nigel's advice from 1988 and dubbed the AGI machine in question "Big Mother"). I also included a few nascent ideas for the "outer cognition" part of an AGI, with which I was generally pleased at the time (except for one step of my "Big Mother roadmap", now called C01, which was very hard, and I just didn't have the time to think it through properly, so that section was really crap!) Finally, I wrote a second, autobiographical, part, as a way of saying "I was here, I existed!" that would persist in the cloud forever, long after there was no-one left alive who remembered me.

And so, on Friday 16 March 2018, I went to sleep on a double mattress in the back of my car, barricaded in a lock-up garage in Coventry, with no food or water, and DNR written on my chest in black marker, waiting to die from dehydration, which I expected to take anything from 3-10 days.

The next day, after 36 hours in the garage, I changed my mind, and drove home. It wasn't that I'd formulated some magic alternative plan; I still believed homelessness to be the unavoidably most likely outcome, for all the reasons previously stated. But I figured that I'd been the SAS of risk for ~25 years, so why stop now...? (Plus, I missed my cat, who I'd deposited safely in a cattery!)

## Recent history

My landlord was amazing (he let me stay on at Rancho Relaxo rent-free for two months), and Aubrey was amazing, lending me some money to help me tidy up my affairs in Northampton and relocate to Cambridge (considering that I had to move anyway, I figured that I might as well move to the hottest AI/tech hotspot in the country, in order to increase my chances of finding a job).

And so I moved back to Cambridge (and another tiny bedsit) - back to where it all began in 1983.

(I couldn't bring my cat, so I put him into temporary digs, from which he escaped after just 3 days; he's been on the lamb ever since. By now, he's either dead, or shackled up with a little old lady...)

I figured that it was reasonable, under the circumstances, to make use of all the help that was available, so I applied for Jobseeker's Allowance, and Housing Benefit, and enrolled with a couple of organisations that help people get back into work. I even joined a food bank - every little helps!

To be honest, in my heart I still believed that all I was doing was delaying the inevitable; for almost a year, I must have thought about suicide as a longer-term possibility roughly every week. I would see events advertised for 2 or 3 months in the future, and question whether I would still be alive by then. And I certainly didn't expect to ever see how Game of Thrones season 8 turned out!

I was systematic in my search for work, doing everything that you're supposed to, leaving no stone unturned. I would take advice, adopt a strategy, focus on that strategy, and then change up if it hadn't worked out after a couple of months. But nevertheless the reaction from employers (and there are thousands of AI and high-tech jobs in Cambridge) was exactly as I'd predicted. But I persevered regardless, and Aubrey was just fantastic, ultimately PayPal-ing me money from California every week for over a year (totalling god knows how many £000s!) to keep me going.

Being new in town, I joined [meetup.com](https://www.meetup.com) to find out about events etc that were happening in Cambridge, and joined a few AI-related groups. This then led to the following sequence of ideas:

- starting my own AGI-themed Cambridge-area meetup group
- recruiting Cambridge-area volunteers to collaborate on progressing Big Mother to completion
- starting a SENS-like non-profit to continue with the MMM AI project, only on a non-profit basis

Unsure, I asked Aubs what he thought about the latter, to which his reply was "absolutely, go for it!" And so, having failed to progress the MMM AI project on a for-profit basis, and still being technically alive (and therefore unable to let go of my lifelong dream), I decided to transform the MMM project into a "non-profit". Given what I had already been through, what did I have to lose?

And so, in August 2018, I founded bigmother.ai cic (<https://bigmother.ai>), a community interest company, and also the CambridgeAGI meetup group (<https://meetup.com/CambridgeAGI>).

This would obviously now have to be my "weekend hobby", rather than my day job (assuming that I was ultimately successful in finding traditional paid programming work), but nevertheless:

- it was a way of keeping the AGI project going, rather than simply giving up on it entirely
- having "non-profit AGI volunteer work" on my CV might act as a catalyst in my search for work
- the long-term possibility now existed that e.g. Bill & Melinda Gates might PayPal the Big Mother project £10 million or whatever, thereby allowing the project to hire me as a full-time employee

Either way, having the AGI project back in my life gave me something to live for that I loved doing.

I decided to completely review my description of Big Mother as it appeared in “How to build an intelligent machine” (with which I wasn’t entirely happy), and to then deliver the revised and improved version as a series of CambridgeAGI talks, which I would video and put on YouTube. This was actually a lot of work, especially considering that most of my time was spent job hunting.

Speaking of which, by December 2018 I had applied for about 100 jobs, and had endured a number of interviews (as well as “programming exercises”, which is a big thing these days), every single one of which ended in a soul-destroying rejection just as I had anticipated back in January.

After one particularly disastrous interview experience, I was really starting to think that I’d never find work, and that Big Mother might ultimately prove to be the only way out of my predicament.

- 2019 -

By March 2019, having been thinking about it in the back of my mind for about a year, I finally worked out how to do Big Mother roadmap step C01 properly, at least in outline. Prior to this the roadmap comprised 19 steps, each nominally associated with a future workgroup (techie team), and I could see in my mind’s eye a path to completing 18 of them, but not the 19th, not C01, and without that last piece I couldn’t see a solution to the AGI problem as a whole. And so C01 was the final piece of a 35-year jigsaw puzzle, which might not sound like much but for me it was a massive step. Just as Aubs had seen a path to his verifier, and “all” that was required was the resources (and stamina) necessary for him to complete it (to the point at which he seems to have had a ~25 year lead on the state of the art, no less), I could now see a path to completing each of the 19 individual roadmap steps required to build a maximally-safe, maximally-benevolent AGI, which means that (from a very high-level project management perspective) “all” that was now required was the resources (and stamina) necessary for 19 coordinated teams to complete their individual parts, and a maximally-safe, maximally-benevolent AGI would pop out at the other end.

This was my second “I know how to do it!” moment, and “all” I needed now was to put 19 expert teams together and raise the necessary \$5-10 billion funding, which is what bigmother.ai is for.

In April, contrary to all prior expectations, I finally found out how Game of Thrones turned out!

Then, 449 days after driving home from that garage in Coventry, I was finally offered a job! :-)

But I couldn’t have done any of it without the help and support of many people, particularly Aubs!

## **Current status**

So, after all that, I now have *two* jobs:

- my “A” job (weekdays) - started Monday 1 July 2019
- my “B” job (evenings and weekends) is to continue to progress the Big Mother project. I reckon it will take 50-100 years to complete safely, so my life’s plan is simply to take it as far as I can.

## **What happens next...?**

In all likelihood:

- after I give my CambridgeAGI talks, the “real” AI community will laugh me out of Cambridge :-)
- ... or they won’t recognise the merit of my ideas until after I’ve been dead for a decade or so!
- without the support of the mainstream AI world, Big Mother will fail to raise significant funding
- without significant funding, Big Mother will fail to make significant further technical progress

Nevertheless, it should now be obvious to even a casual observer that I’ll never give up even if the odds are a billion to one against, and, accordingly, even if some or all of these things do actually occur, which some of them surely will, while I’m still technically alive I will just carry on regardless!

Aaron Turner, September 2019