

THE ROYAL TIMES  
**MAGAZINE**  
15.12.18



**PROBABLY THE SMARTEST  
BRAIN IN BRITAIN  
(HE'S THE TECH GENIUS FROM FINCHLEY  
WITH DESIGNS ON YOUR FUTURE)**

Demis Hassabis is a chess prodigy turned games designer whose company, DeepMind, is leading the race to develop artificial intelligence. His mission? To use it to solve everything in our lives. Tom Whipple meets the publicity-shy geek



Demis Hassabis, 42,  
photographed by Immo Klink



**F**or the CEO of a half-billion-pound company, Demis Hassabis's working day begins very moderately. There is no 5am kale smoothie with *The Wall Street Journal*, no intensive workout followed by breakfast board meeting. Instead, he arrives in his office beside King's Cross in London at a sensible hour, spends the day in meetings, then gets home in time for a family dinner with his wife and two children at 7.30pm.

There he is able to relax and put the kids to bed before, at 10pm, beginning what he calls his "second day of work". And that is when things get a little less moderate.

"I stay up until around 4am," he says. "Sometimes 4.30, depending how it's going." If the first "day" of his working day is about running the business, guiding its 700 employees – 400 with PhDs – as they seek to maintain their position as the world's leading artificial intelligence company, the second day is about reminding himself why he runs this company at all. It is about computer science, mathematics and keeping up to date with the field. "That's when I do my research work."

Although, he admits guiltily, in the past couple of weeks, that "research" has to a large degree involved analysing chess matches.

Hassabis's company, DeepMind, bought by Google in 2014 for £400 million, is the closest thing Britain has to an organisation with Elon Musk-style ambition. It is not aiming to be incremental, but transformative. Its mission: "To solve intelligence and then to use that to solve everything else."

When we meet, Hassabis, 42, is in reflective mood. A few days after our interview, DeepMind will do something rather odd for a supposedly money-making enterprise. It will publish an academic paper, describing a key part of its intellectual property – and a crucial step in achieving that improbable mission statement.

It would not be hyperbolic to say that paper is the final word on the most astonishing story so far in the history of artificial intelligence. It is about a program his company designed called AlphaGo, which was then refined into another program called AlphaZero. It didn't only solve one of the biggest problems in artificial intelligence – defeating humans at the strategy board game Go – it also did so in a way so general that it could, as an afterthought, become by far the best chess computer in the world, too.

It could be argued – and this is something philosophers have no shortage of disputes about – that it is in a genuine way "intelligent". It is also, for Hassabis, the culmination of a journey that took him to Cambridge and University College London, that made him a chess master at 13 and multimillion-selling



Hassabis aged 7, at home in Finchley with one of his chess trophies

**'I'D BE MUCH MORE PESSIMISTIC ABOUT THE WAY THE WORLD IS GOING TO GO IF I DIDN'T KNOW THAT AI IS ON THE WAY'**



Hassabis at 23, running his second games company, 1999



With Stephen Hawking, 2017

software designer before he was an adult, but which began when he was a young child in Finchley, north London, beginning to wonder how the human mind works.

AlphaZero, however, is not the reason that Demis Hassabis has spent his "second working day" of late analysing chess matches. He has done that because some old friends are in town. When we meet, London is hosting the final of the World Chess Championship, which, after 12 tense draws, was eventually won by the defending champion, the Norwegian Magnus Carlsen.

This meant that, amid the travelling circus of trainers, commentators and assorted grandmasters, Hassabis had been reliving his youth. "It's been a festival of chess," he says. "I was joking with Judit Polgar [a famous grandmaster] just now actually. I hadn't seen her for 30 years." He and Polgar are almost exactly the same age and, when they were 12, they were fierce competitors. "She was No 1 when I was No 2." When he says "No 2", he means in the world – Polgar is Hungarian. They were the best in their age group.

That was his peak, though. To the shock of pretty much everyone, he gave up chess. He never made grandmaster level and watched as Polgar went on to be the finest female player in history. Chess, he had decided, wasn't for him. "I love what these people do. There's an incredible artistic purity to putting your whole brain and mind on it," he says. "I knew early on that I was much better as a generalist."

For a former child prodigy with a ferociously intimidating intellect and a superhuman work schedule, Hassabis does not even have the good grace to be socially deficient. In person, while his short stature and glasses make him look a little like a suave Penfold from *Danger Mouse*, there are no weird tics, no missed eye contact. He is not only normal for a computer scientist, depressingly he seems normal for a normal person.

In a way, that's why he quit chess. "There are too many fascinating things in the world to spend your whole life obsessing," he says. "That's how I felt when I was young. I loved physics, art, lots of things. To become Magnus Carlsen you have to be a severe obsessive. I am obsessive, but in a general way – if that's not an oxymoron."

Hassabis's CV appears, at first, to be one of an extremely high-achieving dilettante. He flits from discipline to discipline, from business to university to university to business, without alighting anywhere – at least, until now. But perhaps the most astonishing thing about it (and you are about to see that that is no mild claim) is that it was planned.

The most logical place to begin the story of his career is – it really is – when he was four. That was when he first became intrigued ➤

by a chessboard. “Chess is where it all started for me. Chess is about improving decision-making. Until I was 12 or 13 I considered myself a professional chess player. Albeit a kid. You’re studying really hard, talking to trainers, analysing your own games, thinking about how you can improve.

“It got me into thinking about the process of thought: what is intelligence, how is my brain coming up with these ideas?” So, just like that, he quit.

Or at least, he quit top-flight chess. David Silver, a fellow programmer and computer science academic, first met him a few years later, when they were both 15. “I played on the junior chess scene in Suffolk,” says Silver. “Demis was the kid who would turn up from London when he needed more pocket money. He’d win the local tournament, take the £50 prize money – that was how he earned his pocket money at the weekends.” These days Silver is a programmer at DeepMind, where he led the AlphaZero project.

Much of that prize money went on what was fast becoming Hassabis’s new hobby. “I bought my first computer when I was 8, with the winnings from a chess tournament,” he said. “My parents had no idea about computers, still don’t. They didn’t say anything though – it was my money. I started programming. I realised computers were an amazing, almost magical, device for extending the power of your mind.”

You are probably thinking, where were his parents in this? What hothousing was involved? Who was pushing him? And when did he have an inevitable teenage breakdown?

The answer is, they weren’t really involved at all. His Greek Cypriot father was a toy salesman, his Chinese-Singaporean mother worked at John Lewis. He describes them as “quite bohemian”, and says that they always encouraged him and his two siblings to pursue their interests. For his brother and sister that meant music and creative writing; for him it meant something they just did not understand. “My parents lost track of what the hell I was doing when I was 14 or 15.”

That was also when he moved out. Between finishing his A levels early and going to university there was a hiatus. He had been accepted to Cambridge, but they said they wouldn’t have him until he was older. So he entered a programming competition with the games company Bullfrog, and came second.

“They invited me in for a summer job.” At that point, they learnt – like Cambridge – that he was under-age. “They said, ‘You can’t work here. It’s illegal to employ you.’ I said, ‘Can I come in anyway?’ I was paid £200 a week in brown paper envelopes. I paid for my YMCA hostel room in cash.”

Here, he gained far more than a summer job. He co-developed *Theme Park*, one of the most successful games of the Nineties,



The DeepMind offices at Google’s King’s Cross HQ



## AT THE AGE OF 16 HE WORKED ON A RUDIMENTARY AI. ‘I KNEW THIS WOULD BE THE GREATEST ADVENTURE YOU COULD TAKE’

which shifted 15 million copies. In particular, he worked on a rudimentary AI, which meant it responded to how people played it. He was 16 and he began to see his future. “When I saw how successful that was, I realised this would be the greatest adventure you could take.” From then on, “It was a pretty coordinated plan; I always had a company like DeepMind in mind.”

Like a chess master thinking 15 moves ahead, he just had to position the pieces, making his different career and academic gambits until everything was in the position to strike.

The first piece to be positioned was Cambridge University, where an undergraduate computing degree augmented his practical programming skills with more theoretical ones. “It was closer to maths than computer science. I loved all that. That gave me a good grounding.” Although he disagreed with their approach to artificial intelligence. “It was all about logic. It was obvious to me it wouldn’t get anywhere. You’d never be able to program in enough knowledge for it to suddenly become intelligent. We don’t know enough about knowledge in our own minds to encapsulate it as a bunch of rules.”

The second move, after Cambridge, was a return to making computer games – this time

with him in charge. The goal, though, was not programming but business. “I needed to learn about running a company and teams.”

After releasing two games – aficionados may remember *Black & White* and *Republic: the Revolution* – he had learnt enough and closed his games studio. But, “It still didn’t feel time to start DeepMind. I needed a source of ideas other than computer science. That’s when I did my PhD in cognitive neuroscience. Specifically, in areas of neuroscience we had little idea how to adapt to artificial intelligence.” He wanted to learn how the brain worked, so that he could replicate it in a computer.

At last, in 2010, it was time.

DeepMind, which Hassabis cofounded with friends Mustafa Suleyman and Shane Legg, is still run as a separate company. Ever since Google bought it – its biggest European acquisition – DeepMind has been based in the company’s UK headquarters in King’s Cross.

The office is everything you would hope for from a technology giant. It is the kind of place where you are rarely more than a few feet from a ping-pong table – and if you are, that’s probably because you are beside the smoothie machine. At the front desk the pierced and studiously dressed-down receptionists give off the air that this is just a temporary gig, before their interpretive dance careers take off. The windows are decorated with retro, pixellated, video games characters.

The only thing that slightly lets down the ubercool image is the employees themselves, who walk past with T-shirts tucked in, invariably carrying laptops, exuding that whiff, unmistakable once you know it, of a childhood spent in Games Workshop. It is a little like the Cambridge University computer science course has suddenly been transported, blinking, to a hip Dalston coffee shop. Which is, of course, not that far from the truth.

But you could take away all the glass, steel and artfully unfinished MDF walls, and you would still be left with something else. There are as many Go boards as ping-pong tables, and the equation-filled whiteboards in the halls get more use than the table football machines in front of them. On a tour I pass the Turing office, the Von Neumann office and the Ramanujan conference room. “There’s been a big push for women,” says one employee as we pass the two Marys – Shelley and Somerville.

Throughout there is a tremendous, almost tangible, sense of intellectual excitement. It feels fizzing, like there could be anything behind the doors. Somehow it is not surprising to go into the Gauss room and find a grandmaster and international master watching the tie-break matches of the World Chess Championship, while using AlphaZero to analyse. “It just plays exquisite games,” ➤



says Matthew Sadler, the grandmaster in question, as AlphaZero recommends a move missed by Magnus Carlsen.

Then I shut that door and wonder what is behind the next one.

To understand why AlphaZero's games are so special, why Sadler is so excited, you have to go back 21 years. In 1997 a computer called Deep Blue beat Garry Kasparov, the world chess champion, in one of the most famous matches ever played. It had been billed as the final battle of man versus machine.

Actually, though, there was another game that computers could not beat: Go, a Far Eastern board game with many more playing positions than there are atoms in the universe.

Go is not a game that responds to brute-force calculation. It requires intuition and an instinctive appreciation of positions and beauty. Unlike chess, where by 2016 a mobile phone could play a credible game against even a grandmaster, there were no programs that came close to playing at the top level of Go. Many thought at that point that Go had another 20 years of human supremacy left. Hassabis thought differently.

Google bought DeepMind because it wanted it to advance artificial general intelligence, to design programs that could not just be good at one thing – such as chess – but many. Eventually, such a general intelligence could be used to solve scientific and commercial problems.

DeepMind has, in fact, already done some of these things – its software is used to recommend purchases on Google Play, and has slashed energy consumption in Google's server rooms. Two weeks ago it won a competition to predict how proteins fold – an apparently esoteric skill that is actually crucial to all processes in the body. But a good way of proving its intelligence to the world was to do so in the same way that humans – including Hassabis – prove their intelligence, using games.

The approach DeepMind took to Go was very different from programs that had gone before. While Deep Blue required a team of programmers and grandmasters to teach it the tricks of chess, to impart the canon of human knowledge accrued over millennia, AlphaZero required nothing – literally. This program started life like a baby born into a world of Go. All it knew was that it wanted to win, and all it had was a board and a knowledge of the rules.

Then, like a baby learning its first steps, over the course of tens of thousands of games played against itself, it experimented and refined its understanding. It learnt – for want of a better word – intuition. Slowly, it independently gathered the skills of human masters, as well as a few more of its own. Then it used those skills to demolish a redoubt of human intellect. The infancy of AlphaZero was measured in hours – at

Playing chess with Mickey Adams, Chris Flowers and, far right, world champion Magnus Carlsen



## IN THE MEDIUM TERM, AI COULD LEARN TO WORK LIKE A SCIENTIST AND 'MAKE A MAJOR BREAKTHROUGH AT THE NOBEL PRIZE LEVEL'

the end of which it was able to defeat the best human in the world.

That was not enough though. After his program became the undisputed Go computer champion Hassabis went to a conference, to describe the achievement. Murray Campbell, who had been on the Deep Blue team that defeated Kasparov, approached him. "Murray asked, 'What would it do if we tried it on chess?'" For a conventional program, the question would be meaningless – Deep Blue played grandmaster chess, but asking it to play Go would be as relevant as asking a pigeon to. It would require starting again and designing a new system from the ground up. For AlphaZero, it just required a different infancy.

So it was that a new AlphaZero was born, again into a world of black and white. This time though, the board held not the flat counters of Go, but the knights, kings, bishops and queens of chess. "AlphaZero could start in the morning playing completely randomly, then by tea would be superhuman level," said Hassabis. "By dinner it would be the strongest entity there's ever been." When he looked through the games, he realised it was not just the best player in the world; it was not just better than the best computer; it had also discovered a new way of playing.

"It's like chess from another planet," said Hassabis. "You can appreciate this beauty. It's sacrificing pieces, opening up new attacking lines." One grandmaster, allowed to analyse its games, compared it to finding the old notebooks of a long-forgotten chess genius of the past.

Between the table football and the soaring ambition, all of this might, justifiably, make DeepMind sound like a parody of Silicon

Valley. Except, if it is, it is Silicon Valley before the fall. Hassabis is the computer scientist's computer scientist. In this programming Valhalla he has achieved something that eluded Facebook, Microsoft and all those before him: he has kept collisions with the tedium of analogue reality to a minimum.

Not for him monopolies commissions or Senate hearings about Russian election interference. When DeepMind was bought out by Google, it didn't just gain cash; it gained freedom.

"There was a big decision point as to whether we should sell or not sell the company. There's a lot of trade-offs on both sides," he says. In selling, he lost one kind of autonomy – he suddenly had a boss. That is not the only autonomy there is, though. "One thing was weighing on my mind. It never looked attractive to me to be a publicly quoted CEO. Then your life is just earnings calls and quarterly targets, not pure fundamental research. I prefer science."

Google was clear this purchase was not like its one of YouTube. DeepMind is not expected to turn a rapid profit. It is there to solve the problem of intelligence; this is a Manhattan Project of computing. "I realised if I did this with Google I could be behind the scenes and get on with my science," says Hassabis.

There is still a lot to do. Intelligence is not solved simply by winning a few board games. Chess is an immensely complex problem, but also an immensely simple one. Unlike life, it has clear rules, clear outcomes and everyone knows what is going on. There are real-life situations analogous to this – ones where, like chess or Go, the solution is about choosing the right path through lots of different possible combinations. One they are already making progress on is understanding protein-folding, crucial for designing new drugs. Our inability to model protein chemistry in the body is one reason developing new drugs costs so much.

One of the weirdest things about artificial intelligence programs is that humans create them but do not understand them. Once they are set running, learning and developing, their workings can become as opaque to outsiders as our own brains.

For Hassabis, the chess program is a way of seeing AlphaZero's brain. On a chessboard, unlike in proteins, you can see code made manifest. "I can viscerally feel the optimisation process happening on a board," he says. By solving a problem he understands – chess – he can gain an insight into one he doesn't, such as proteins. "If the knight is not optimal, AlphaZero might move it into six positions to get it right. If you are good enough at chess it's almost like you can feel it. It's like being in the mind of AlphaZero."

Hassabis, naturally, has bigger plans than merely saving millions in pharmaceutical ➤

research. He thinks that, in the medium term, AI could learn to work like a scientist, devising hypotheses and designing experiments to test them, and then “make a major breakthrough at the Nobel prize level of challenge”.

Occasionally, between dreams of Stockholm and analysing chess games, reality does rudely interfere. One of the company’s projects involves analysing health data. The same pattern recognition that can intuit the best response to a late Queen-side castling can also guess the best diagnosis of, say, an early-appearing retinal condition. To be able to make that diagnosis though, DeepMind needs a lot of data – it needs access to NHS patient records.

This has been where people have started taking notice. In 2017, the Information Commissioner censured the Royal Free Hospital for not doing enough to inform patients about the partnership before it handed over data to DeepMind for them to process on its behalf. Just last month, it was announced that the company’s health subsidiary will transfer to Google – leading to calls for greater assurances that UK patient data will be ringfenced.

These problems are inevitably going to keep on intruding, wheedling their way into even the rarefied atmosphere of DeepMind’s

## **‘NOW I’M WORLD CHESS CHAMPION BY PROXY,’ HE SAYS. ‘THERE WAS PROBABLY SOME UNFINISHED BUSINESS’**

Plato room (“the cave” – basically like the other rooms, but with beanbags). Most will be dealt with – some by the company itself, some by society adjusting to a new world of big data and algorithms.

The one that seems unlikely ever to go away is also the most fundamental. Instead of wondering whether we can develop true artificial intelligence, some argue we must start worrying whether we should.

Hassabis, like everyone in the field, has a rehearsed answer to what could broadly be characterised as the “robot apocalypse problem”. He positions himself as cautiously concerned but not hyperbolic. There are ethical problems with artificial intelligence, but they can be dealt with, and the difficulties it may solve are more of a draw than the ones it might cause.

“I’d be much more pessimistic about the way the world is going to go if I didn’t know there was something as game-changing as AI on the way,” he says. “There are so many problems out there, from Alzheimer’s disease to climate change, that are hugely

complex and where we seem to be making almost no progress. Either we need an exponential improvement in human behaviour or an exponential improvement in technology, and the world doesn’t look like it’s getting its act together on the former.”

For now, though, he is not focusing on the apocalypse. Instead, he is enjoying an achievement that is as much about the past as it is the future. At the age of 13, the young Demis Hassabis made a calculation. He gave up chess, despite knowing that he had a chance to be the best in the world.

Then he took a path that, he admits, seems quite random. Yet like the knight that made six moves just so it could get in the right place to strike, he has moved towards something that was, somehow, inevitable all along.

To today, when in a room in his office, staffed by extremely excited grandmasters, he has created the best chess player in the world. “I’ve come full circle. I’m world champion by proxy. I guess if you’re going to be Freudian about it,” he says, “there was probably some unfinished business.” Checkmate. ■