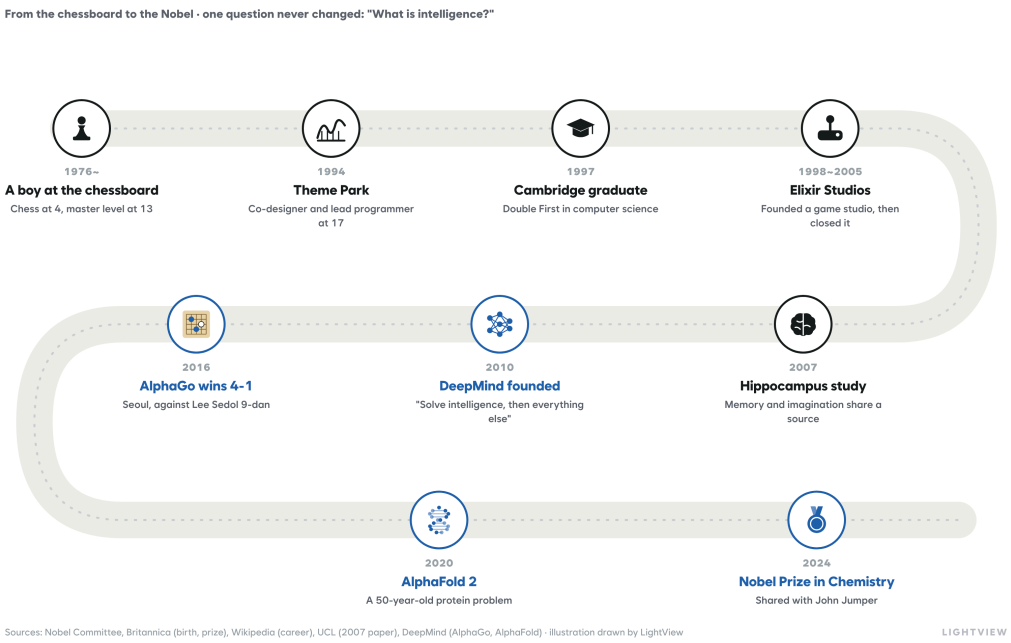
DEEP DIVE LIGHTVIEW DEEP DIVE · PEOPLE

# From the Chessboard to the Nobel: Why Hassabis Chased Intelligence

Forty years of Demis Hassabis, the man who has led Google's AI

**Part 1** The past · 1976–2024

Part 2 · Google, OpenAI and Anthropic, now and next




## Key points

1. **It starts with a 12-year-old's question.** After resigning a 10-hour chess game, the boy asked whether this was really the best use of so many brilliant minds.
2. **Games and neuroscience were two roads to one question.** The developer who built small worlds in Theme Park went on to show, in his PhD, that memory and imagination rely on the same part of the brain, the hippocampus.
3. **DeepMind proved, through games, that machines can teach themselves.** From Atari to a 4-1 win over Lee Sedol to AlphaGo Zero, which learned with no human games at all.
4. **The goal was science from the start.** AlphaFold cracked the 50-year-old protein-folding problem, and Hassabis shared the 2024 Nobel Prize in Chemistry.

**LIGHTVIEW DEEP DIVE** · A series that looks closely at the people and companies leading industry and technology. Based on public sources and reporting. Not investment advice.

# 1 A boy at the chessboard

**Demis Hassabis was born in London in July 1976 to a Greek Cypriot father and a Chinese Singaporean mother. He learned the rules of chess at four by watching adults play, and by 13 he had reached master level (Elo 2,300), ranked second in the world for his age.**

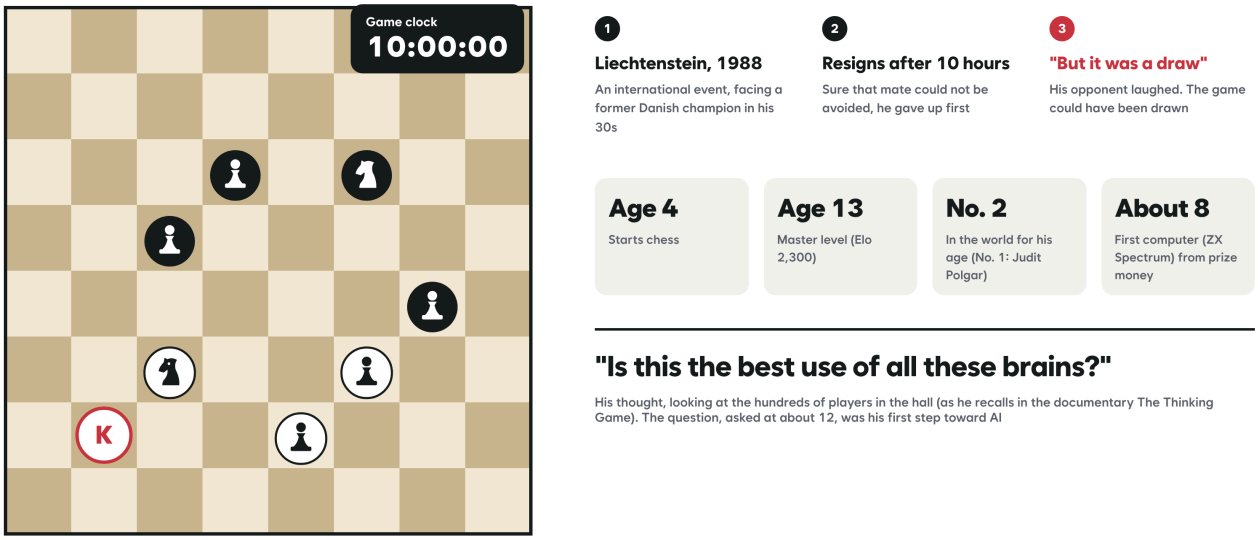



**Figure 1 · The ten-hour game** The moment a chess prodigy began to look beyond the board. The game is illustrated, not reproduced

The turning point came at an international tournament in Liechtenstein in 1988. Ten hours into a game against a former Danish champion in his thirties, the boy decided checkmate could not be avoided and resigned. His opponent asked why he was giving up when the position was a draw. Hassabis recalls feeling sick to his stomach.

Looking at the hundreds of players filling the hall that day, he began to wonder whether spending such brilliant minds on chess was the best use of them. What drew him was less chess itself than the thinking that produced the moves. The programming he had taught himself on his first computer, a ZX Spectrum bought with chess prize money when he was about eight, became the tool for that question.

At around 11, he recalls, he beat his younger brother at Othello with a program he had written himself. Winning because his own program had done the thinking for him was his first hands-on taste of the idea of a thinking machine.

He remained strong enough at chess to captain England's junior teams. But the question he held onto for life was not how to win. It was where winning thoughts come from.

> **"I was more interested in the actual thought processes going on behind the games I was playing"**
>
> Hassabis, talk at the University of Cambridge (reported by Entrepreneur, Mar 26 2025), paraphrased

# 2 Building small worlds in games

**After chess, his next stage was video games. A game is a small virtual world with rules and goals, and making its characters behave believably is an artificial intelligence problem.**

Told by Cambridge to take a year off first, he joined the studio Bullfrog. At 17 he co-designed Theme Park (1994) with the famed Peter Molyneux and was its lead programmer. Visitors in the game react to every ride and food stall, and it sold millions. Bullfrog offered a large sum to keep him; he went to university.

After a top-honours computer science degree in 1997, he led AI on Lionhead's Black & White, whose creature learns by watching the player, and in 1998 founded his own studio, Elixir. It never found the success he hoped for and closed in April 2005.

The failed founder became a student again. Most AI then, game AI included, was rules written by hand, which he felt could never capture the complexity of the world. The only working general intelligence was the human brain, so he set out to study it.

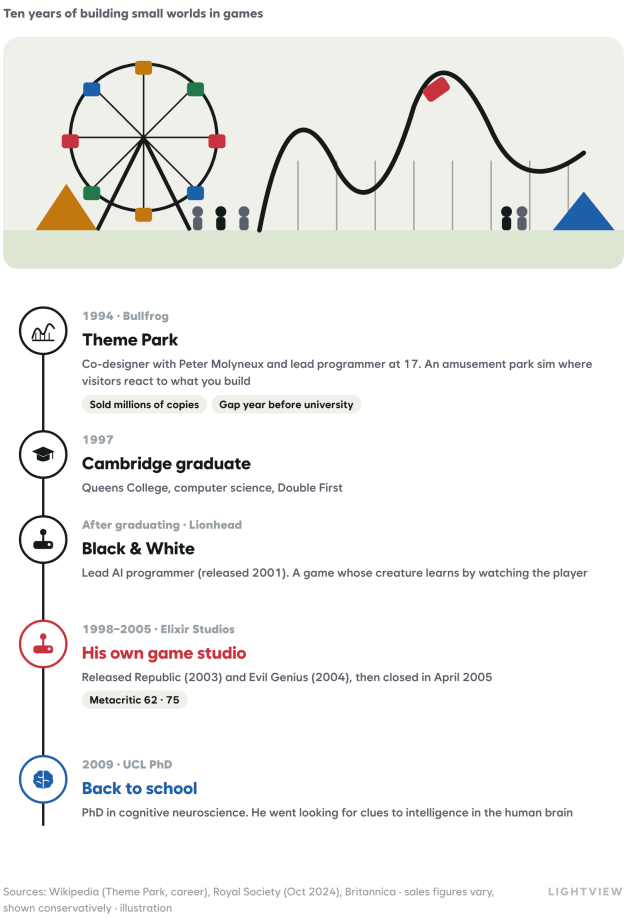



**Figure 2 · Ten years in games** From Theme Park to Elixir, then back to school

### What games taught him, and what carried into DeepMind

| Game years | Lesson | Carried into DeepMind |
|---|---|---|
| **Theme Park** | Set the rules, and a different story plays out every time | Training AI inside virtual worlds |
| **Black & White** | A creature that learns by watching the player | Reinforcement learning from rewards |
| **Elixir, founded and closed** | The gap between an ambitious goal and the market | Why he chose Google as a backer |

The third column is LightView's reading of his career and later choices

Theme Park had no fixed storyline. The developers set only the rules for visitors and rides, and each player's park played out differently. It taught him first-hand the power of a small world run by rules: a simulation.

That instinct is still with him 30 years later. Google's "world models", covered in Part 2, aim to let AI simulate the world in its head: the game developer's small worlds, carried into the study of intelligence.

# 3 Memory and imagination share a source

Hassabis studied cognitive neuroscience at University College London (UCL), earning his PhD in 2009. He focused on two things machines were worst at, memory and imagination, and picked the hippocampus, deep inside the brain, as the key.

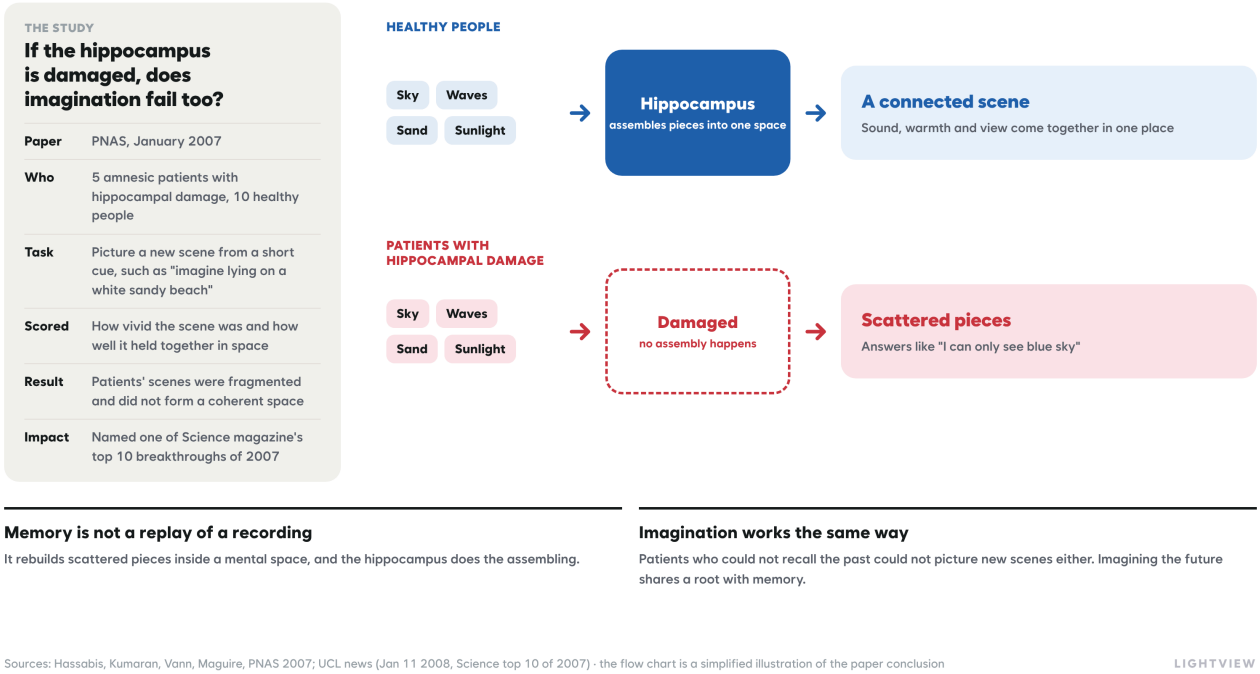



**Figure 3 · The hippocampus and imagination** Patients with hippocampal damage could not recall the past, and could not picture new scenes either

In a paper in PNAS in January 2007, the team gave five amnesic patients with hippocampal damage and ten healthy people tasks such as "imagine you are lying on a white sandy beach". Healthy people described one connected scene of sun, waves and sand; the patients' scenes fell apart into fragments.

The conclusion was that memory is not rewinding a tape but reassembling pieces in a mental space, and that imagination does the same thing to picture scenes that have not happened yet. Science magazine named the study one of the top ten breakthroughs of 2007.

> ## Real intelligence means building a world in your head and simulating the future inside it
>
> The idea that took shape here would lead, two decades later, to Google's world-model strategy (Part 2)

### Hints from the brain, carried into AI

| What the brain does | What AI borrowed | Where |
|---|---|---|
| **The hippocampus replays the day's scenes to strengthen memory** | "Experience replay": store past game frames and learn from them again | DQN (2015, ch. 4) |
| **It assembles pieces into one space to picture new scenes** | "World models": AI pictures the world and tests ideas inside it first | Genie and others (Part 2) |

The DQN paper (Nature 2015) says experience replay was inspired by replay in the hippocampus

# 4 A game with no rules given

| 2010 | 3 | Mission |
|---|---|---|
| DeepMind founded in London | Hassabis, Shane Legg, Mustafa Suleyman | "Solve intelligence, then use it to solve everything else" |

At UCL, where he was doing postdoctoral research, he met Shane Legg. Together with Mustafa Suleyman, a friend of his younger brother whom he had known since childhood, they founded DeepMind in 2010. Their goal was not a program that does one job well but a general intelligence that can learn anything by itself.

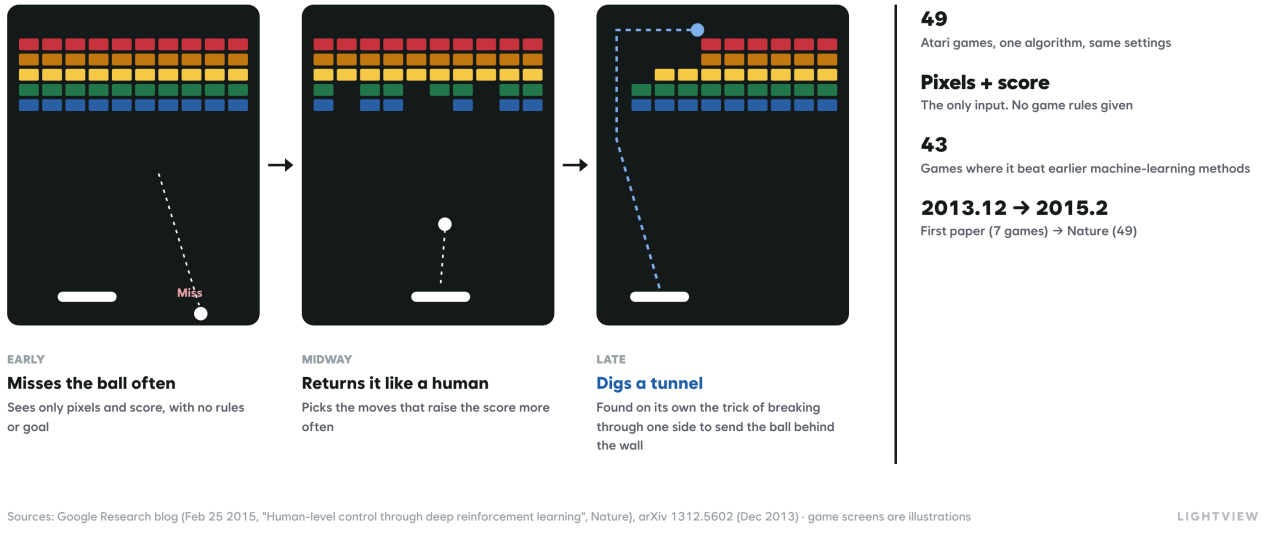



**Figure 4 · Atari Breakout** An AI that knew no rules found, by trial and error, a tunnel strategy few people use

The first training ground was video games: the result is immediate, and you can try again endlessly. DQN, first described with seven games in December 2013, appeared in Nature in February 2015, learning 49 Atari games the same way from screen pixels and score alone and beating earlier methods in 43. In Breakout, it even worked out on its own how to break through one side and send the ball behind the wall.

### DeepMind's first four years

| When | What happened |
|---|---|
| **2010** | Founded in London. Early backers included Founders Fund (Peter Thiel), Horizons Ventures, Elon Musk and Jaan Tallinn |
| **Dec 2013** | First draft of the DQN paper, an AI that teaches itself Atari games |
| **Jan 2014** | Acquired by Google (reported at $400–600m). An ethics board was reportedly a condition |
| **Feb 2015** | DQN published in Nature, beating earlier methods in 43 of 49 games |

# 5 Choosing Google, proving it in Seoul

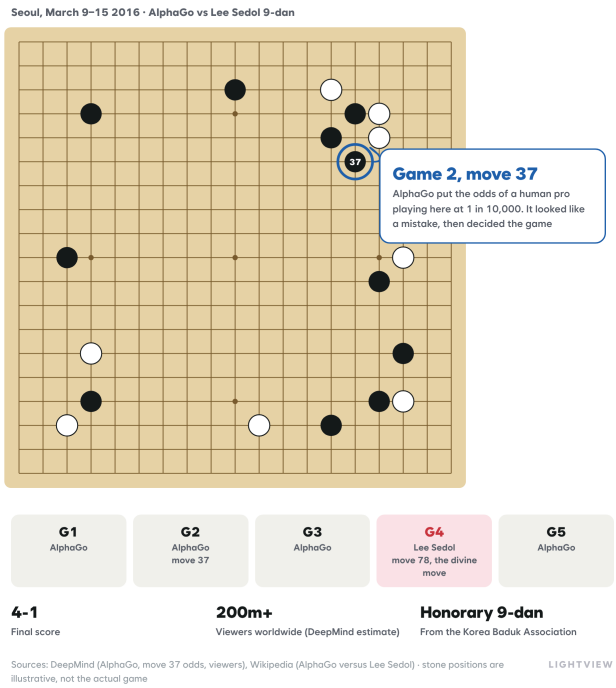



**Figure 5 · AlphaGo vs Lee Sedol** It won four of five games. Stone positions are illustrative

Google bought DeepMind in January 2014. The price was never announced but was reported at $400–600m, and DeepMind had reportedly turned down Facebook. Some reports said setting up an internal ethics board was a condition of the deal.

Hassabis has said it would have taken ten years to build Google's computing infrastructure on his own. Being able to focus on research, instead of spending his time raising money and building a business as a startup CEO, is also seen as a reason he chose Google.

Two years later the choice paid off in Seoul. In March 2016 AlphaGo beat Lee Sedol 9-dan 4-1. Move 37 of game two was one AlphaGo itself estimated a human pro would play with a probability of 1 in 10,000. Lee took one game with move 78 of game four, and more than 200 million people watched worldwide.

What mattered was not just the win. AlphaGo learned the basics from human games, then improved by playing countless games against itself, and along the way found moves people had not thought of.

**AlphaGo's matches**

| Opponent | When | Result | Note |
|---|---|---|---|
| **Fan Hui 2-dan (European champion)** | Oct 2015 | 5-0 | First computer to beat a pro |
| **Lee Sedol 9-dan** | Mar 2016 | 4-1 | Seoul, 200m+ viewers |
| **Top online players ("Master")** | Dec 2016–Jan 2017 | 60-0 | Played online under a hidden identity |
| **Ke Jie 9-dan** | May 2017 | 3-0 | Wuzhen, China, then world No. 1 |

Sources: DeepMind, Wikipedia (AlphaGo versus Lee Sedol, Ke Jie, Master) · retiring in November 2019, Lee Sedol said there is "an entity that cannot be defeated", even if he were number one

There is a reason Go became AI's test. The number of possible board positions is about 10 to the 170th power, far more than the number of atoms in the observable universe (about 10 to the 80th). Reading moves broadly and deeply, as chess programs do, is not enough. It takes something closer to intuition.

AlphaGo combined one neural network that picks the next move with another that judges the position, and looked deep only into promising moves. It was a machine that learned how to judge, not just how to calculate, and a step toward the general intelligence Hassabis had talked about.

# 6 Throwing away human games

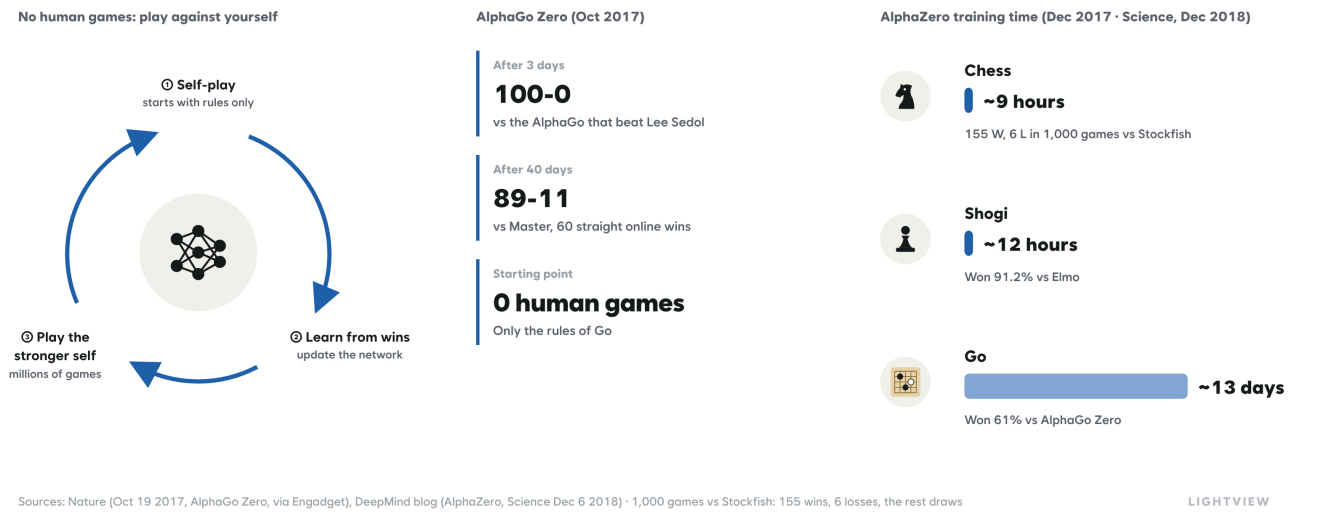



**Figure 6 · The self-play loop** Given only the rules and left to play itself, it went beyond human knowledge

DeepMind went one step further. AlphaGo Zero, unveiled in October 2017, started from the rules of Go alone without seeing a single human game. After three days of teaching itself, it beat the version that defeated Lee Sedol 100-0, and after 40 days it beat "Master", winner of 60 straight online games, 89-11.

AlphaZero, unveiled that December, learned chess, shogi and Go with a single method. In the 2018 Science paper, after about nine hours of chess training it played 1,000 games against Stockfish, the strongest chess engine, winning 155 and losing 6. Nearly 30 years after that game in Liechtenstein, an AI he built stood at the top of chess.

**DeepMind AI, sharpened on games**

| Name | Year | What it did | Why it mattered |
|---|---|---|---|
| **DQN** | 2013–2015 | Learned 49 Atari games from pixels and score | One method, many tasks |
| **AlphaGo** | 2016 | Beat Lee Sedol 9-dan 4-1 | Found moves people had missed |
| **AlphaGo Zero** | 2017 | No human games; beat AlphaGo 100-0 after 3 days | Surpassed humans without human knowledge |
| **AlphaZero** | 2017–2018 | Chess, shogi and Go with one method | Generality |
| **AlphaStar** | 2019 | Top 0.2% in StarCraft II (Grandmaster) | Real time, hidden information |

Sources: Nature (Feb 2015, Oct 2017), Science (Dec 2018), Nature (Oct 2019, AlphaStar), DeepMind blog

Games had a clear limit, though. Their rules are exact, wins and losses come immediately, and you can replay them millions of times. Real scientific problems are not like that. One experiment can take months, and the answer cannot be checked right away. So Hassabis took the methods honed on games and went looking for a scientific problem whose answer could be checked.

# 7 Solving a 50-year-old problem

**Games were practice. From the start Hassabis was aiming at science's hardest problems, and the first target was protein structure prediction. A protein is a chain of amino acids, and the 3D shape that chain folds into determines what it does in the body.**

Shape matters because of medicine. If you know the shape of a disease-causing protein, you can find where a drug molecule could attach, and designing new enzymes or vaccines gets faster. The trouble was how hard shapes were to work out. Ever since Christian Anfinsen, winner of the 1972 Nobel Prize in Chemistry, showed that the amino-acid sequence determines the shape, scientists had tried to compute shapes from sequences alone, and kept failing. Shapes had to be solved one at a time by experiment, sometimes taking years for a single protein.

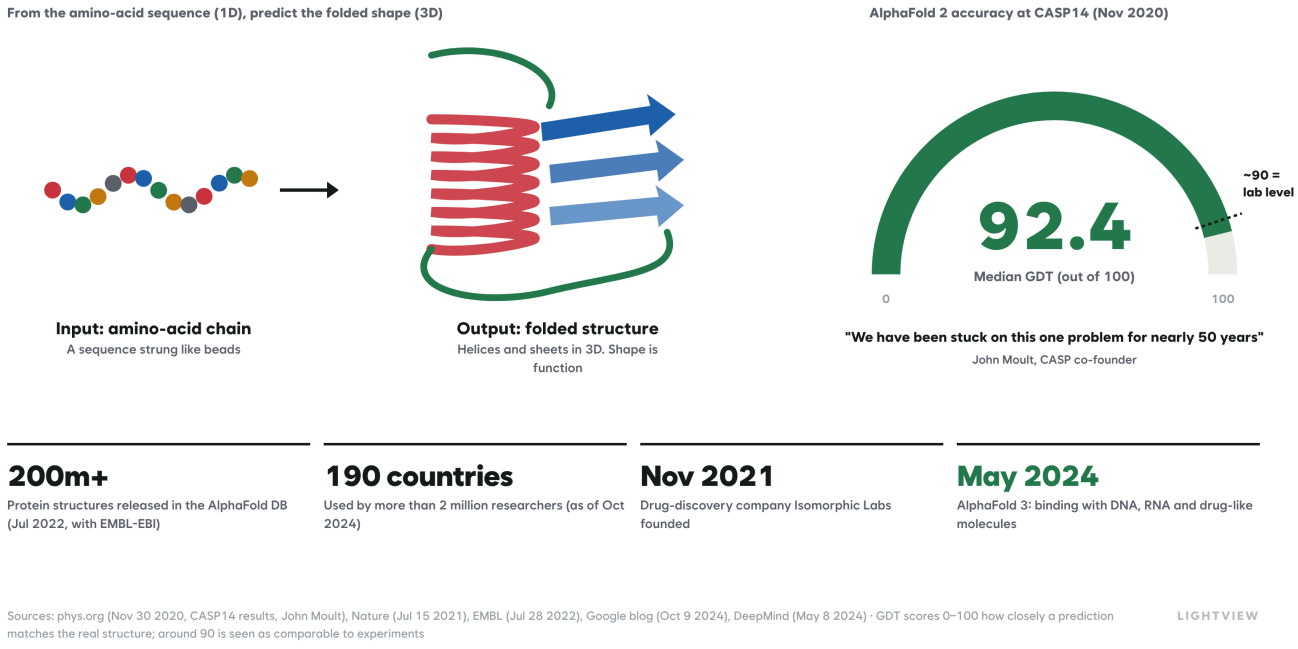



**Figure 7 · AlphaFold** Predicting a 3D structure from a 1D sequence. In 2020 it reached accuracy comparable to experiments

AlphaFold took first place in 2018 at CASP, the biennial competition for this problem, and in November 2020 AlphaFold 2 scored a median of 92.4. A score of around 90 is considered comparable to structures obtained by experiment. The reaction of the scientist who started the competition shows how much it meant.

> **"We have been stuck on this one problem, how do proteins fold up, for nearly 50 years"**
>
> John Moult, CASP co-founder (Nov 2020, reported by phys.org)

DeepMind did not keep the result to itself. It made more than 200 million predicted structures free for anyone to use, and set up a separate company, Isomorphic Labs, to pursue drug discovery directly. DeepMind highlights cases where these structures have been used in research such as malaria vaccines and plastic-eating enzymes.

AlphaFold has not solved everything, of course. How proteins move and change shape, or how many molecules interact in complex settings, remain hard, and predictions still need to be checked by experiment. Even so, by cutting a first step that used to take years down to a computer calculation, it is widely seen as changing how science is done.

# 8 The Nobel, and the question left open

**On October 9 2024 the Nobel Prize in Chemistry went half to David Baker and half jointly to Hassabis and John Jumper. The committee said AlphaFold had fulfilled "a 50-year-old dream".**

The chessboard boy's question had reached science's highest honour some 30 years on. He had been knighted for services to AI that March, and he hoped AlphaFold would be remembered as the first proof of AI's power to speed up discovery.

Yet his Google was struggling on a very different battlefield. OpenAI's ChatGPT, launched in November 2022, had moved the centre of AI from the lab to consumers' phones.

> **Will "solve intelligence first" work in the chatbot race too?**
>
> Part 2: the ChatGPT shock, Google DeepMind's comeback, his 2026 step back as CEO, and Google vs OpenAI and Anthropic

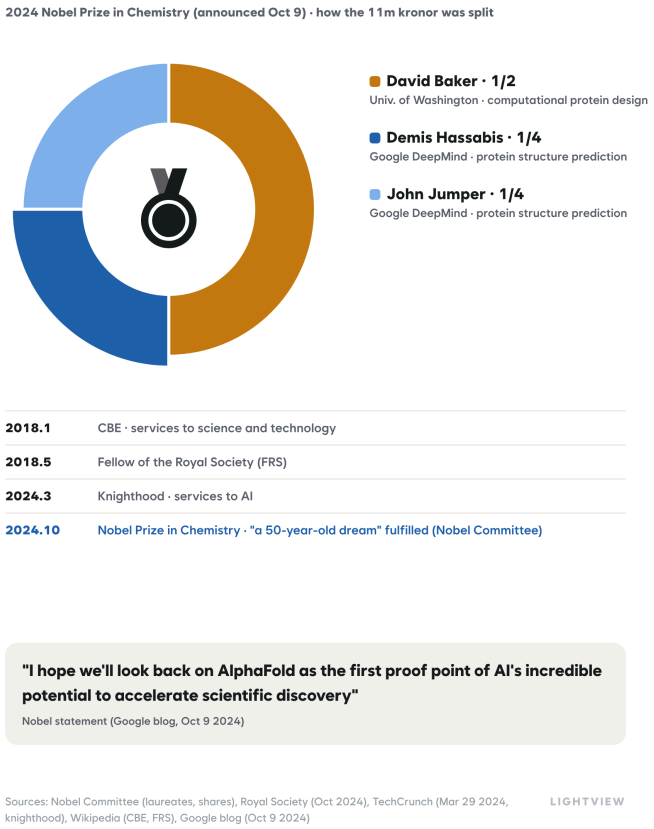



**Figure 8 · The 2024 Nobel Prize in Chemistry** How the prize was shared, and Hassabis's main honours

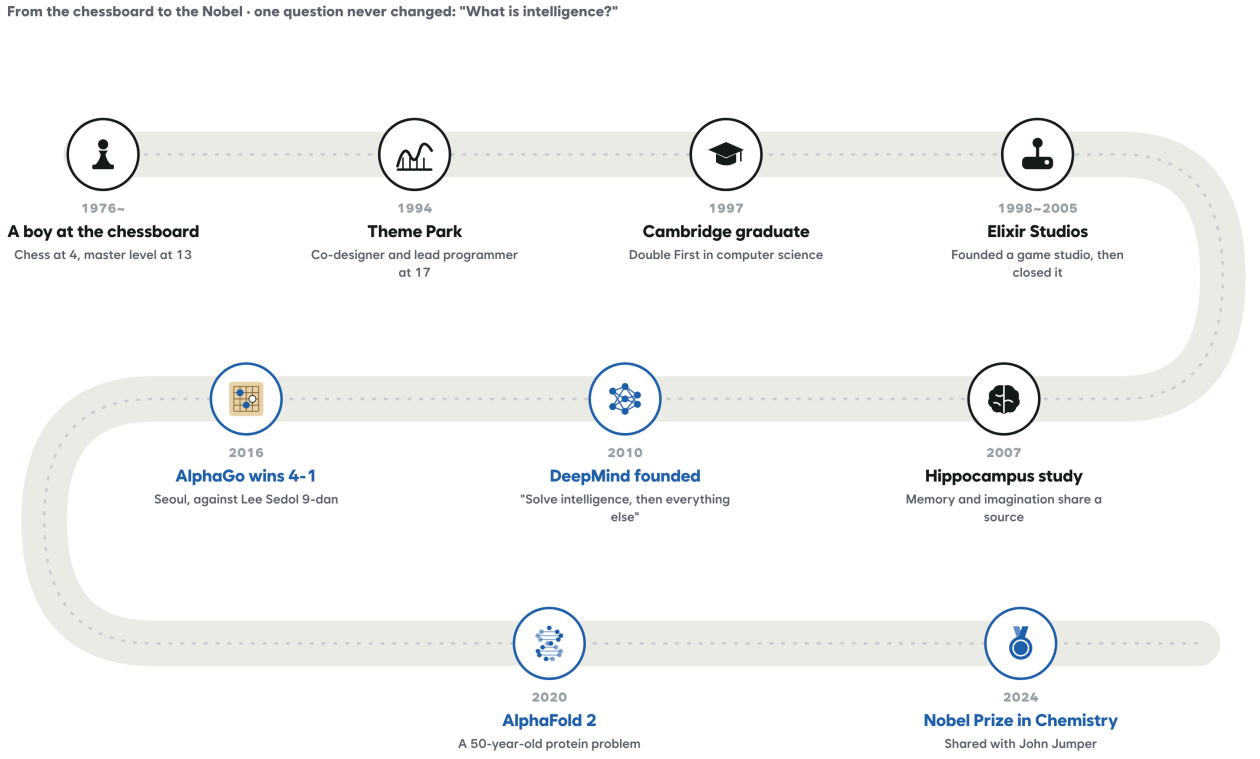



**Figure 9 · Forty years at a glance** One question throughout: what is intelligence, and what can it solve?

# Terms and sources

## Terms

**Artificial general intelligence (AGI)** — AI that can handle work across many fields like a person, rather than one task.

**Reinforcement learning** — Learning that improves behaviour through trial, error and reward (score).

**Self-play** — An AI learning by playing against itself.

**CASP** — A biennial protein structure prediction competition that uses newly solved structures as test questions.

**Hippocampus** — A region deep in the brain, central to forming and recalling memories.

**Neural network** — A computing structure modelled on connections between brain cells; the basic framework of today's AI.

**Protein structure prediction** — Working out the 3D shape a protein folds into from its amino-acid sequence alone.

**GDT** — A 0–100 score of how closely a predicted structure matches the real one.

### Early years and games

Nobel Committee laureate facts (birth) · Britannica (Demis Hassabis) · ChessBase (Oct 9 2024, age ranking) · Google DeepMind documentary The Thinking Game (summarized by OfficeChai, Nov 29 2025, Liechtenstein game) · Wikipedia (first computer, Theme Park, Elixir) · Royal Society (Oct 2024, Theme Park at 17)

### Neuroscience

Hassabis, Kumaran, Vann, Maguire, "Patients with hippocampal amnesia cannot imagine new experiences", PNAS (Jan 2007) · UCL news (Jan 11 2008, Science top 10 of 2007) · Mnih et al., Nature (2015, experience replay)

### DeepMind and game AI

The Guardian / The Observer (Feb 2016, mission) · City AM (Jan 27 2014), Mail & Guardian (Jan 28 2014, acquisition, investors, ethics board) · Backchannel (2015, Google infrastructure quote) · Google Research blog (Feb 25 2015, DQN) · DeepMind (AlphaGo, move 37, viewers) · Nature (Jan 2016 AlphaGo · Oct 19 2017 AlphaGo Zero) · Science (Dec 6 2018, AlphaZero) · Nature (2019, AlphaStar)

### AlphaFold and the Nobel

phys.org (Nov 30 2020, CASP14) · Nature (Jul 15 2021) · EMBL (Jul 28 2022, 200m structures) · DeepMind (May 8 2024, AlphaFold 3) · DeepMind AlphaFold impact stories (malaria vaccine, plastic-degrading enzymes) · VentureBeat (Oct 9 2024, prize statement) · Google blog (Oct 9 2024) · Royal Society (Oct 2024) · TechCrunch (Mar 29 2024, knighthood) · Nobel Prize (1972, Christian Anfinsen)

The outline of this piece drew on a published video commentary; every number and fact was checked directly against the sources above. It reflects information public as of November 8 2026. Drawings are used instead of photos of Hassabis himself.

---

The editor chooses the topic and angle and reviews and revises each column before publishing. AI tools were used for research, drafting and the infographics, and every number was checked against its source. The AI tool used (Claude) is made by Anthropic, one of the companies compared in Part 2.