DeepMind

Research careers, scientific collaborations, and a London laboratory’s changing place inside Google

DeepMind is an artificial-intelligence research organization founded in London in 2010 by Demis Hassabis, Shane Legg and Mustafa Suleyman. It became widely known through AlphaGo, which defeated Go champion Lee Sedol in 2016, and AlphaFold, a system for predicting the three-dimensional structures of proteins. Google acquired the company in 2014. In April 2023, DeepMind and Google Research’s Brain team were combined as Google DeepMind, bringing the London laboratory into the same organization as another major tradition of machine-learning research and engineering.[1, 2]

The name now covers Gemini model and application development alongside frontier research. Its leadership also changed in August 2026: Hassabis became chair of Google DeepMind and chief scientist of Alphabet, while Koray Kavukcuoglu became senior vice president of Google DeepMind, reporting to Google and Alphabet chief executive Sundar Pichai. Kavukcuoglu oversees Gemini model development, frontier AI research, and the Gemini app and developer teams; he also remains Google’s chief AI architect.[3, 4]

Within this history are some unusual working lives. A researcher in speech recognition turns toward ancient inscriptions. A professor of robot learning works on magnetic control in a fusion experiment. Philosophers write about how the benefits of advanced AI should be distributed. These are real activities, but their coexistence does not mean that employees can simply choose anything and keep doing it. To understand the latitude, one has to follow how an interest acquires collaborators, an experiment, a sponsor and enough continuing resources to become someone’s work.

London and Google, 2010–2021: founding and financial support

Hassabis brought experience from computer games as well as academic science. He had founded the games company Elixir before studying imagination and memory with neuroscientist Eleanor Maguire at University College London. His subsequent research crossed UCL, MIT and Harvard. Legg’s route passed through mathematical work on intelligence at the Dalle Molle Institute for Artificial Intelligence in Switzerland and postdoctoral research at UCL’s Gatsby Computational Neuroscience Unit. The institution’s combination of learning algorithms, cognition and ambitious engineering was present in its founders’ formation; it was not added later as a diversity of applications.[5, 6]

The early work nevertheless needed specific tests. In their December 2013 Atari paper, Volodymyr Mnih, Kavukcuoglu, David Silver, Alex Graves and colleagues described a network that learned actions from screen pixels and reward signals. The same architecture and learning algorithm were tested on seven games. The problem joined visual representation to learning through action: a system had to work out which patterns mattered for future rewards rather than receive a handcrafted account of each game.[7]

This gave the organization a productive relation between a large ambition and a small experiment. Games supplied repeatable environments, scores and relatively cheap trials. The experiment did not demonstrate general intelligence. It did make a broad research direction susceptible to comparison, debugging and collective improvement. A researcher could pursue a question larger than a product feature while still being answerable to evidence.

Google supplied a different kind of experimental capacity. In a 2015 interview with Steven Levy, Hassabis said that reproducing Google’s computing and engineering infrastructure independently would have taken about a decade. He also emphasized Larry Page’s substantive interest in AI. In his account, the acquisition provided both an unusually sympathetic patron and machinery that could shorten the cycle between an idea and a result.[8]

[External image / 外部圖片] The entrance to Google’s building at 6 Pancras Square in London, photographed in 2017.
The entrance at 6 Pancras Square in December 2017, when Google and DeepMind occupied the building. This is a historical location, not a photograph of the organization’s 2026 headquarters. Photograph: Gciriani, Wikimedia Commons, CC BY-SA 4.0. Full composition retained.
The financial arrangement is visible in a later, carefully bounded document. DeepMind Technologies Limited’s 2021 accounts recorded £1.365 billion of turnover as remuneration for research and development from other companies in the group, against £1.254 billion of administrative expenses. The accounts identify infrastructure, staff costs and stock-based compensation among the reasons expenses increased. They also record a written assurance of financial support from Google LLC. These are figures for a UK legal entity in 2021, not the present worldwide Google DeepMind budget. They show why calling DeepMind “profitable” without explaining internal group payments would miss the underlying arrangement: Alphabet was paying for a research capability.[9]

The arrangement moved much of the judgment about research expenditure to the parent company and to managers allocating staff and computing capacity.

AlphaGo and learning through self-play, 2015–2018

David Silver’s group carried the learning-and-control program into Go. AlphaGo combined networks that proposed moves and assessed positions with a search through possible

continuations. Training used human games and then games against versions of the system itself. It defeated European champion Fan Hui in 2015 and Lee Sedol by four games to one in Seoul in March 2016. Go expertise supplied both a demanding opponent and a community capable of judging the play.[10]

The subsequent work changed the experiment. AlphaGo Zero dispensed with the human-game examples; AlphaZero extended learning from self-play to chess and shogi as well as Go. Silver, Thomas Hubert, Julian Schrittwieser and colleagues reported the fuller evaluation in 2018. The aim was no longer simply to beat another Go player, but to investigate how much of the machinery could carry across games with different rules. The public match was one stage in a continuing technical program.[11]

Researchers, engineers and the organization of experiments

Neuroscience remained a working subject rather than merely a founding metaphor. In 2018, Jane Wang, Zeb Kurth-Nelson, Matt Botvinick and colleagues used recurrent neural networks to investigate how a system could learn the structure of a task and adapt quickly to new examples. They recreated six experiments from neuroscience and compared the model's behavior with earlier experimental findings. The work proposed an account of relationships among dopamine, prefrontal cortex and learning. Its output was a scientific hypothesis, not an immediately saleable application.[12]

Joel Leibo's work offers another continuity beneath an apparently wandering portfolio. His current research page connects computational neuroscience, cultural evolution, cooperation and institutional economics with human-compatible AI. In 2018 he helped introduce Psychlab, a virtual laboratory in which humans and agents could perform comparable psychological tasks. Rather than treating a good video-game score as a complete measure of intelligence, the researchers built tests of particular capacities. Later language-based social simulations offered a different instrument for studying interaction. The continuity lies in the questions about cognition and cooperation, not in a single product category.[13, 14]

Such work requires people whose contribution is not chiefly proposing the next theory. In a 2022 company interview, research engineer Akhil Raju described mornings spent repairing failures from previous experiments and preparing robots, followed by coding, meetings and impromptu whiteboard discussions. His route had led from robotics competitions through MIT, a startup and Google. He had assumed that a doctorate might be necessary to enter DeepMind; a recruiter instead recognized the relevance of his robotics experience. This is one dated account, not a representative timetable, but it puts experimental upkeep back into the picture.[15]

Aliya Rysbek's work, described in another 2022 interview, was further from the public demonstration. As a full-stack engineer on the Platform team, she helped build an internal system showing research projects, participants and connections among teams. Design,

deployment, documentation, testing and bug fixing made the organization's own activity legible to its members. Her path from Kyrgyzstan through scholarships in Turkey and Hungary, mentoring and more than seventy job rejections also complicates the idea of an effortless entrance into this world. These were rejections across organizations, not seventy applications to DeepMind.[16]

Research latitude is partly made out of this work. An idea is easier to pursue when someone has maintained the environment, made experiments reproducible, or made it possible to find the colleague with the needed expertise. The visible scientist's freedom and the engineer's obligations belong to the same arrangement.

AlphaFold, 2016–2026: an experiment becomes a program, then disperses

Hugo Penedones has described an early route into protein research that is more revealing than the eventual prize. In a Portuguese interview with INESC TEC, he recalled a three-day internal hackathon in February or March 2016. Rich Evans proposed protein folding; Evans, Marek Barwiński and Penedones tried reinforcement-learning and optimization experiments interacting with the FoldIt game. An internal prize and enthusiasm from leadership helped turn the experiment into full-time work. Penedones distinguishes his participation, ending in 2017, from the later AlphaFold2 achievement.[17]

A short interval and an internal demonstration had supplied a route to sponsorship. Leadership could then commit time to a problem whose solution remained uncertain.

John Jumper brought a different preparation. After leaving an initial physics doctorate and working on computational models of proteins, he completed a chemistry PhD at the University of Chicago. Work in the joint orbit of Karl Freed and Tobin Sosnick exposed him to theory and experiment, including a brief laboratory rotation making proteins. He later emphasized the value of knowing what experimentalists find easy, difficult or uncertain. He joined DeepMind in 2017. This was interdisciplinary research assembled from substantial prior training, not a talented generalist casually taking up biology.[18]

Quanta's account of the project describes a team that expanded from six to fifteen after Jumper's arrival. Its first success in CASP, the Critical Assessment of protein Structure Prediction competition, did not secure an indefinite license to refine the same approach. Leadership wanted a more decisive solution or a different important problem. The team rebuilt. The 2019 intern Raphael Townshend recalled the force of concentrated expertise and computing resources. Quanta separately describes the team testing predictions against biological work beyond a competition score.[19]

AlphaFold2's 2021 paper reported a redesigned system incorporating biological and physical knowledge, information from related sequences, and training data accumulated by experimental structural biologists. Its major result concerned predicting structure from

amino-acid sequence. It was not a complete account of the physical process of folding, every changing molecular state, or all of biology. The distinction matters institutionally: an ambitious program had a tractable target against which progress could be judged.[20]

[External image / 外部圖片] A scientific figure comparing AlphaFold2 predictions with experimental protein structures and outlining the model architecture.
Evidence and construction in the same figure: CASP14 results, comparisons of predicted and experimental structures, and the AlphaFold2 architecture. John Jumper and colleagues, 2021, Nature; image via Wikimedia Commons, CC BY 4.0. Uncropped.
Making the result useful demanded another institution. The European Bioinformatics Institute at EMBL helped turn predictions into a searchable public database. The initial 2021 release expanded in 2022 to more than 200 million predicted structures. Scientists did not all have to reproduce DeepMind's computational work to benefit from it. EMBL-EBI's expertise in maintaining biological data, backed by an intergovernmental organization, was part of the accomplishment.[21]

[External image / 外部圖片] Rows of equipment in an EMBL-EBI data centre, with a person working in the background.
An EMBL-EBI data centre, photographed in February 2026. The institute's data-maintenance role is part of the wider infrastructure behind public biological resources; this image does not identify a particular AlphaFold server. Photograph: Jeff Dowling / EMBL-EBI, Wikimedia Commons, CC BY-SA 4.0. Full composition retained.
Release was itself a research and governance task. DeepMind's account of the AlphaFold decision describes consultation with more than thirty outside experts and a multidisciplinary committee reviewing projects, partnerships and publications. One question was how to communicate uncertainty rather than simply suppress less-certain predictions. This is evidence of the 2022 process, not proof that identical procedures or principles remain in force today.[22]

A separate commercial route had opened in 2021, when Hassabis announced Isomorphic Labs as a new Alphabet company pursuing AI-based drug discovery. DeepMind and Isomorphic subsequently developed AlphaFold3 together. Its 2024 launch extended prediction to interactions among proteins and other molecules, while its public server was offered for noncommercial research. A freely searchable structure database, a research model and a drug-discovery company are different ways of organizing the value of the same scientific lineage.[23, 24]

By 2026, the dedicated team had been dissolved, according to Scientific American's reporting following the Financial Times. Some researchers moved within Google or to Isomorphic Labs; others left. Jumper announced in June that he was leaving for Anthropic. A company spokesperson told Scientific American that many moves had happened more than a year earlier. The database and prediction server remained available. Scientists interviewed for the article distinguished completing a well-defined mission from exhausting the field's unanswered questions.[25]

This is a different ending from either failure or permanent protection. A successful program can leave behind a public scientific resource, commercial descendants and dispersed people without retaining its original team. The available reporting does not let us assign a precise proportion of that decision to scientific completion, commercial pressure or organizational restructuring.

Inscriptions and fusion: collaboration beyond the laboratory

Yannis Assael’s epigraphy work began through a personal connection outside his immediate field. In an Oxford alumni interview, he recalled returning to the university during his early DeepMind years and discussing ancient history over lunch with a friend doing doctoral research. He had worked on speech and lip-reading systems. The historical problem opened another use for sequence modeling: damaged inscriptions whose missing text, origin or date could be difficult to establish.[26]

The resulting work did not consist of an AI researcher supplying answers to passive historians. Ithaca, published in 2022, joined DeepMind with researchers in humanities, classics and informatics at Ca’ Foscari, Oxford and Athens. Assael, historian Thea Sommerschield and their collaborators developed tools for restoration and geographical and chronological attribution. Historical context remained part of the task, rather than an inconvenience to be eliminated by a model.[27]

[External image / 外部圖片] Broken marble fragments carrying an ancient Greek inscription, assembled to show surviving text and substantial gaps.
The Hekatompedon inscription, IG I^3 4B (EM 6794), concerning the Athenian Acropolis, in the Epigraphical Museum. This photograph makes the material problem of missing text visible; it is not an AI-generated reconstruction. Photograph: Marsyas, May 29, 2007, Wikimedia Commons, CC BY-SA 3.0. Full composition retained.
Aeneas extended the collaboration to Latin inscriptions in 2025. Its roughly 176,000-inscription corpus drew on existing epigraphic databases; twenty-three historians participated in evaluation. The system retrieved contextual parallels as well as assisting restoration, with an interface and open research materials intended for use beyond the original team. This was a sustained specialty, supported by scholarly labor accumulated elsewhere, rather than an isolated curiosity briefly tolerated by a technology company.[28]

Assael’s current CV lists a Google DeepMind career beginning in 2017 and a senior staff research-scientist role from 2026. It places AI for culture alongside audiovisual speech, multi-agent communication and foundation-model evaluation and post-training. The combination is striking, but a list of interests does not reveal how his time or computing allocation is divided. What it establishes is that an unusual line of inquiry can coexist with more central model work in one research career.[29]

Fusion research entered through another institution’s problem. At a London hackathon in

2018, EPFL researcher Federico Felici explained the difficulty of controlling the magnetic fields that confine plasma in a tokamak. EPFL brought a simulator developed over more than twenty years, physical expertise and access to its TCV experimental device. DeepMind brought reinforcement-learning methods and an interest in controllers that could learn across demanding settings. Policies trained in simulation were subsequently tested on the machine.[30]

[External image / 外部圖片] A general view of EPFL’s TCV tokamak installation, surrounded by experimental equipment and access platforms.
EPFL’s TCV installation in an archival photograph uploaded in 2006, before the DeepMind collaboration. The physical apparatus and accumulated control expertise were essential inputs to the later research. Credit: CRPP-EPFL, Association Suisse-Euratom, via Wikimedia Commons, CC BY-SA 2.5. Full composition retained.
Martin Riedmiller, whose Freiburg professorship and work on robot learning preceded DeepMind, described the practical limits in a 2023 interview. A small number of trials on the tokamak might be followed by a week’s wait. Simulation fidelity therefore mattered enormously. The team also chose to learn control rather than merely imitate an existing controller. A wide research remit could coexist with scarce experimental access and a tightly specified technical choice.[31]

These collaborations help explain why the portfolio can look more miscellaneous than the underlying work. A sequence model, a learned controller or an experimental environment can travel between disciplines. But travel takes translation: historians must explain what counts as a plausible parallel; plasma physicists must explain what a simulation leaves out. Neither inscriptions nor fusion becomes an inexpensive playground simply because machine learning can be applied to it. The fusion result addressed control; commercially viable fusion remained a separate challenge.

Choosing a direction, and deciding to change it

The organization’s own AI-for-science account describes a combination of proposals from researchers and priorities set by leadership. It stresses consequential questions, suitable data and methods, and meaningful ways to evaluate progress. It also distinguishes exploration from the more concentrated engineering required once a promising direction emerges. Research engineers and program managers help organize that transition. This is a stated operating philosophy, not a disclosed approval chart or evidence that every project receives equal consideration.[32]

Neel Nanda’s account supplies a useful comparison at team level. Interviewed in July 2025, he described substantial autonomy in his interpretability and safety work, including room to anticipate problems before they became immediate production needs. Yet persuading a production team to adopt a method was a separate difficulty. He explicitly warned that autonomy and management vary considerably among teams. His own move into leadership had been contingent: he joined expecting to be an individual contributor

and became a lead after his predecessor stepped down.[33]

One can therefore have freedom to investigate without authority to decide how a deployed model will be built. The permissions are related but distinct. This also makes the individual manager and group unusually important: joining the organization is not the same as joining the particular working arrangement described in a favorable interview.

In a December 2025 talk, Nanda explained a change in his team's direction. Much of its 2024 effort had gone into sparse autoencoders, tools intended to disentangle patterns inside models. Improving reconstruction metrics for months had not supplied enough evidence of practical usefulness. The team shifted toward tasks connected to safety goals and comparisons with simpler methods. Nanda did not declare the older approach worthless; he retained a place for discovering unexpected hypotheses. He questioned whether a convenient technical proxy deserved to remain the main objective.[34]

The public record is much thinner on ideas that never gained sponsorship or quietly disappeared. Successful collaborations and unusually candid talks cannot supply a census of the laboratory's possibilities.

Google Brain, Gemini and the reorganizations of 2023–2026

Google Brain had a distinct history before the merger. A 2012 paper by Quoc Le, Marc'Aurelio Ranzato, Rajat Monga and colleagues, including Jeff Dean and Andrew Ng, trained a large network on unlabeled images using a thousand machines. Its well-known cat-face detectors were part of a broader investigation of learned representations at scale. The institution's advantage included the ability to turn a research question into a distributed-computing experiment.[35]

The 2017 Transformer paper, by Ashish Vaswani, Noam Shazeer, Niki Parmar and five colleagues, proposed an attention-based architecture and tested it on translation. Greater parallelism and lower training time were part of the contribution. This was not a historical DeepMind project. Folding every earlier Google AI result into the DeepMind name would erase the separate teams, engineering traditions and collaborators brought together in 2023.[36]

Pichai's April 2023 merger announcement explicitly connected the combined team to Google's computing resources and future products. It did not merge all of Google Research into DeepMind: other research areas remained in Google Research under James Manyika. The reorganization concentrated responsibility for increasingly capable general models while retaining a wider research system around it.[37]

The August 2026 changes made the balance more explicit. Hassabis relinquished day-to-day operations while retaining strategic roles and leadership of Isomorphic Labs. The announcement also said Jeff Dean and Sanjay Ghemawat would launch an independent public-benefit corporation, with Google as a founding investor and cloud partner. Some

senior researchers’ next institutional arrangement was now outside the organization they had helped build.[38]

Reuters’ August reporting, based on interviews with seven people, described uncertainty among employees accustomed to leaders who had protected projects outside Gemini’s main priorities. It also reported competition for scarce computing capacity on Google’s Tensor Processing Units (TPUs) and changes in reporting arrangements for nontechnical teams. Google declined to answer the detailed questions cited in the report; executives had told staff that day-to-day work would remain the same. These accounts establish tension and uncertainty, not that all exploratory research had stopped.[39]

An official September interview with Kavukcuoglu described parallel research efforts that could run for a year or longer before being brought into a model. Long-horizon work and product pressure are not mutually exclusive. They can be coupled through an expectation that several exploratory tracks will eventually contribute to a competitive system.[40] For a researcher, the important question becomes how far a project may diverge, for how long, and who can defend its continuation when a shared resource becomes scarce.

Philosophers, public debate and the limits of institutional freedom

The DeepMind Institute, launched on September 16, 2026, adds another form of intellectual activity. Its stated purpose is interdisciplinary discussion of artificial general intelligence (AGI) and its social implications. Shane Legg is managing editor; Legg, Manyika and Hassabis are directors. The launch describes a platform involving people from Google DeepMind, Google and the wider research community, and states that essays express their authors’ views rather than Google’s official position.[41]

Some of those authors are employed researchers. The Institute identifies Iason Gabriel as a philosopher and research scientist leading the AGI & Society team, and Atoosa Kasirzadeh as a staff research scientist whose work crosses philosophy, computation and mathematics. Gabriel’s earlier work included the United Nations and Oxford political philosophy. Their joint essay argues for global benefit from AI through considerations of rights, reciprocity, fairness and beneficence. It treats contributions from data, education and public infrastructure as part of the case for benefits extending beyond those who can pay.[42]

Other contributors have external institutional homes. Stephen Cave’s essay on utopianism identifies him with Cambridge and its Leverhulme Centre for the Future of Intelligence and Institute for Technology and Humanity. His argument favors plural, revisable efforts to improve institutions rather than a single final blueprint. Publication on DeepMind’s platform is not evidence that he has become a DeepMind employee.[43]

The platform makes philosophical writing by internal and external researchers publicly available around a shared concern. Its public pages do not disclose fellowships, salaries,

guaranteed research time, or a general selection procedure. Publication alone establishes none of these; nor does advocacy for equitable distribution establish control over commercial decisions.

The limits have appeared in other forms of work. DeepMind's Streams project involved clinicians, engineers and translational researchers building tools for hospital care; the team's transfer to Google was announced in 2018 and confirmed in 2019. The UK Information Commissioner's Office criticized the Royal Free hospital's lack of transparency in sharing records of 1.6 million patients for the project. A research organization's opportunity to experiment can depend on people who did not choose to participate in its institutional ambition.[44, 45]

Employees also distinguish freedom to choose interesting problems from freedom to refuse particular uses. In May 2026 the Communication Workers Union reported that its DeepMind members had voted to seek union recognition and were demanding moral-refusal rights and independent ethics oversight. Management rejected voluntary recognition while agreeing to discussions through ACAS, according to the union. This describes the union's account at that date, not a finding that all staff shared its position or that recognition had subsequently been achieved.[46]

Research careers after the merger

DeepMind's history contains several distinct arrangements for exploratory work: the short experiment that attracted sponsorship, the team assembled around a scientific target, the continuing collaboration with outside scholars, and the philosophical publication platform. Assael's epigraphy work and Nanda's account of changing methods show researchers shaping a direction from within. The engineers' interviews show the maintenance and coordination those directions require.

These arrangements now sit within an organization responsible for Gemini's competitive performance. The dedicated AlphaFold team has dispersed; the Institute has begun publishing; other research careers continue. Whether the next unusual proposal receives sustained support will depend on particular colleagues, managers and resources. The public record shows routes by which that support has been secured, but no general entitlement to it.

Dates and evidence

This profile was researched through October 24, 2026. It distinguishes the original London DeepMind, the pre-merger Google Brain team, and Google DeepMind after April 2023. Employee interviews from 2022–2025 are dated testimony, not claims about every current team. Public sources do not establish a universal protected-time percentage, the present internal budget for each research group, or a complete record of rejected and terminated projects. The current leadership account follows the explicit August 2026 announcement

rather than older biographies still describing Hassabis as chief executive.

---

1. Demis Hassabis, “From unlikely start-up to major scientific organisation”, December 5, 2019; September 2023 curriculum vitae, for historical milestones.↩
2. Sundar Pichai, “Google DeepMind: Bringing together two world-class AI teams”, April 20, 2023.↩
3. Sundar Pichai and Demis Hassabis, “The next chapter of our AI momentum”; current Koray Kavukcuoglu biography, checked October 24, 2026.↩
4. Kenrick Cai, Reuters, “Google shakes up AI leadership as DeepMind chief shifts role”, August 5, 2026. This dates the transition; the official memo supplies the precise titles and remit.↩
5. Hassabis, September 2023 curriculum vitae, pp. 1–2.↩
6. Shane Legg, James Manyika and Demis Hassabis, “Introducing the DeepMind Institute”, September 16, 2026, including author biographies and editorial disclaimer.↩
7. Volodymyr Mnih et al., “Playing Atari with Deep Reinforcement Learning”, December 19, 2013. The seven-game experiment here should not be confused with the later 49-game study.↩
8. Steven Levy, “The Deep Mind of Demis Hassabis”, January 16, 2015. The infrastructure account is Hassabis’s retrospective explanation.↩
9. DeepMind Technologies Limited, 2021 financial statements, pp. 2–4, filed in 2022; archival copy of the company accounts. These are not consolidated 2026 Google DeepMind figures.↩
10. David Silver, “Deep Reinforcement Learning”, June 17, 2016; AlphaGo project account, for the training sequence and matches.↩
11. David Silver and Demis Hassabis, “AlphaGo Zero: Starting from scratch”, October 18, 2017; David Silver et al., “AlphaZero: Shedding new light on chess, shogi, and Go”, December 6, 2018.↩
12. Jane Wang, Zeb Kurth-Nelson and Matt Botvinick, “Prefrontal cortex as a meta-reinforcement learning system”, May 14, 2018, linking the research paper.↩
13. Joel Leibo’s research page, checked October 24, 2026; “Generative agent-based modeling with actions grounded in physical, social, or digital space using Concordia”, December 2023. Simulations are research instruments, not automatic validation of claims about real societies.↩
14. Joel Leibo, “Open-sourcing Psychlab”, January 26, 2018.↩
15. DeepMind, “From LEGO competitions to DeepMind’s robotics lab”, May 19, 2022. Edited employee interview, not independent observation.↩
16. DeepMind, “Kyrgyzstan to King’s Cross: the star baker cooking up code”, May 26,

2022. Role and working practices refer to that interview date.↩

17. INESC TEC, interview with Hugo Penedones on the 2024 Nobel awards, October 16, 2024, in Portuguese. The account used here is Penedones' s firsthand recollection; the article' s incidental CASP chronology is not relied upon.↩

18. Louise Lerner, "How an 'accidental chemist' honed his approach at UChicago on the way to a Nobel Prize", December 3, 2024.↩

19. Quanta Magazine, "How AI Revolutionized Protein Science, but Didn' t End It", June 26, 2024.↩

20. John Jumper et al., "Highly accurate protein structure prediction with AlphaFold", Nature 596, 583–589, July 15, 2021; PubMed record.↩

21. DeepMind, "AlphaFold transforms biology for millions around the world", July 28, 2022; AlphaFold Protein Structure Database.↩

22. Koray Kavukcuoglu et al., "How our principles helped define AlphaFold' s release", September 14, 2022.↩

23. Demis Hassabis, "Introducing Isomorphic Labs", November 2021.↩

24. Google DeepMind, AlphaFold research and release chronology, checked October 24, 2026; the server' s noncommercial-research description concerns the May 2024 launch, not a universal license for every AlphaFold artifact.↩

25. Mary Randolph, "DeepMind broke up the AlphaFold team. Here' s why scientists aren' t alarmed", Scientific American, August 7, 2026. Its statement suggesting research began in 2018 is not used; the 2016 origin is separately sourced above.↩

26. University of Oxford Department of Computer Science, Yannis Assael alumni profile. The page displays December 8, 2021 but contains subsequent updates; it is used as an alumni recollection, not a securely dated snapshot.↩

27. Yannis Assael, Thea Sommerschield, Brendan Shillingford and Nando de Freitas, "Predicting the past with Ithaca", March 9, 2022.↩

28. Google DeepMind, "Aeneas transforms how historians connect the past", July 23, 2025.↩

29. Yannis Assael' s personal CV, checked October 24, 2026.↩

30. EPFL, "EPFL and DeepMind use AI to control plasmas for nuclear fusion", February 16, 2022.↩

31. TalkRL, Martin Riedmiller interview transcript, August 22, 2023.↩

32. Google DeepMind, "A new golden age of discovery: AI for Science", November 2024. An institutional account of research selection and organization, not an independently audited allocation policy.↩

33. 80,000 Hours, Neel Nanda on careers at frontier AI companies, recorded July 21 and published September 15, 2025.↩

34. Neel Nanda, “Our Pivot To Pragmatic Interpretability”, FAR.AI, December 1, 2025.↩

35. Quoc Le et al., “Building high-level features using large scale unsupervised learning”, ICML 2012.↩

36. Ashish Vaswani et al., “Attention Is All You Need”, first submitted June 12, 2017.↩

37. Sundar Pichai, “Google DeepMind: Bringing together two world-class AI teams”, April 20, 2023.↩

38. Sundar Pichai and Demis Hassabis, “The next chapter of our AI momentum”; current Koray Kavukcuoglu biography, checked October 24, 2026.↩

39. Reuters, “Google’s DeepMind shakeup reflects drive to make Gemini an AI leader”, August 13, 2026.↩

40. Google AI: Release Notes, Koray Kavukcuoglu on frontier models, coding agents and building AGI, September 3, 2026, interview and transcript.↩

41. Shane Legg, James Manyika and Demis Hassabis, “Introducing the DeepMind Institute”, September 16, 2026, including author biographies and editorial disclaimer.↩

42. Iason Gabriel and Atoosa Kasirzadeh, “The case for global benefit from AI”, September 16, 2026, including current author biographies.↩

43. Stephen Cave, “Principles for a new utopianism”, September 16, 2026, including author biography.↩

44. Demis Hassabis, Mustafa Suleyman and Dominic King, “Scaling Streams with Google”, November 13, 2018, updated September 18, 2019.↩

45. Information Commissioner’s Office, “Google DeepMind and class action lawsuit”, retrospective account of the 2015 data sharing and 2017 findings. The criticism described here concerns the hospital’s handling of patient information.↩

46. Communication Workers Union, “Google AI Workers to Enter Union Recognition Talks”, May 20, 2026. A union statement about members’ demands and management’s response, not a workforce-wide survey.↩