


Analysis Report


Plagiarism Detection and AI Detection Report
s41586-021-03819-2.pdf

Plagiarism Detection




Plagiarism types	Text coverage	Words
Identical	66.9%	8,081
Minor Changes	24.2%	2,924
Paraphrased	3.5%	422
Excluded		
Omitted Words		0

AI Detection



	Text coverage	Words
AI Text	0%	0
Human Text	100%	12,076
Excluded		
Omitted Words		0

Plagiarism



Results (96)

*Results may not appear because the feature has been disabled.

 Private Cloud Hub 0	 Shared Data Hub 0	 Filtered / Excluded 0
 Internet Sources 78	 AI Source Match 18	 Current Batch 0

Plagiarism types	Text coverage	Words
 Identical	66.9%	8,081
 Minor Changes	24.2%	2,924
 Paraphrased	3.5%	422
Excluded		
 Omitted Words		0

About Plagiarism Detection

Our AI-powered plagiarism scans offer three layers of text similarity detection: Identical, Minor Changes, and Paraphrased. Based on your scan settings we also provide insight on how much of the text you are not scanning for plagiarism (Omitted words).

Identical

One to one exact word matches. [Learn more](#)

Minor Changes

Words that hold nearly the same meaning but have a change to their form (e.g. “large” becomes “largely”). [Learn more](#)

Paraphrased

Different words that hold the same meaning that replace the original content (e.g. 'large' becomes 'big') [Learn more](#)

Omitted Words

The portion of text that is not being scanned for plagiarism based on the scan settings. (e.g. the 'Ignore quotations' setting is enabled and the document is 20% quotations making the omitted words percentage 20%) [Learn more](#)

Copyleaks Shared Data Hub

Our Shared Data Hub is a collection of millions of user-submitted documents that you can utilize as a scan resource and choose whether or not you would like to submit the file you are scanning into the Shared Data Hub. [Learn more](#)

Filtered and Excluded Results

The report will generate a complete list of results. There is always the option to exclude specific results that are not relevant. Note, by unchecking certain results, the similarity percentage may change. [Learn more](#)

Current Batch Results

These are the results displayed from the collection, or batch, of files uploaded for a scan at the same time. [Learn more](#)

🔍 Plagiarism Detection Results: (96)

🌐 Sources – Compute-Structure AI History 92.3%
https://ai-papers.xuzou.cn/sources/2021-alphafold2_2021.txt

Source notes and reports Sources ← Home Article Highly accurate protein structure prediction with AlphaFold <https://doi.org/>...

🌐 Highly accurate protein structure prediction with AlphaFold 64.8%
<https://papers.baulab.info/papers/also/jumper-2021.pdf>

John Jumper

Article Highly accurate protein structure prediction with AlphaFold <https://doi.org/10.1038/s41586-021-03819-2> Received: 11 May 2021 Acc...

🌐 (PDF) Highly accurate protein structure prediction with AlphaFold 63.5%
https://www.researchgate.net/publication/353275939_highly_accurate_protein_structure_prediction_with_alpha...

Home Protein Chemistry General Biochemistry Chemistry Protein Structure ArticlePDF AvailableHighly accurate protein structure predi...

🌐 Highly accurate protein structure prediction with AlphaFold 57.5%
<https://dasher.wustl.edu/bio5357/readings/nature-596-583-21.pdf>

John Jumper

Article Highly accurate protein structure prediction with AlphaFold <https://doi.org/10.1038/s41586-021-03819-2> Received: 11 May 2021 Acc...

🌐 Highly accurate protein structure prediction with AlphaFold | Nature 55%
<https://www.nature.com/articles/s41586-021-03819-2>

Skip to main content Thank you for visiting nature.com. You are using a browser version with limited supp...

🌐 AI Can Now Predict Protein Shapes With Near-Perfect Accuracy | HackerNoon 51.6%
<https://hackernoon.com/ai-can-now-predict-protein-shapes-with-near-perfect-accuracy>

164 reads AI Can Now Predict Protein Shapes With Near-Perfect Accuracy byGoogle byGoogle|@G#8203;googleGoogle develops search, advertising, cl...

🌐 AF2REFPAPERS.md 12.6%
<https://raw.githubusercontent.com/chenxingqiang/alphafold-notebooks/main/alphafold2/af2refpapers.md>

Reference Papers List -----

🌐 Highly accurate protein structure prediction for the human proteome 7.8%
<https://storage.prod.researchhub.com/uploads/papers/2024/02/07/s41586-021-03828-1.pdf>

Kathryn Tunyasuvunakool

Article Highly accurate protein structure prediction for the human proteome <https://doi.org/10.1038/s41586-021-03828-1> Received: 11 May ...

 NPJVACCINES-01030R1rs.pdf https://stacks.cdc.gov/view/cdc/107336/cdc_107336_ds8.pdf rtp1292 Last updated by author(s): Mar 1, 2021 Reporting Summary Nature Research wishes to improve the reproducibility of the work that we publi...	4.9%
 AcroBrwEx_8605_2_supp_88372_sn04j2_ADW858E.tmp https://ora.ox.ac.uk/objects/uuid:c67ad0a7-49e6-40de-b681-c13b847b7239/files/rmc87pr94m jjt2025 Last updated by author(s): Nov 8, 2024 Reporting Summary Nature Portfolio wishes to improve the reproducibility of the work that we publi...	4.7%
 checklist.pdf https://yin.hms.harvard.edu/publications/2019.orbit.sup.sum.pdf aeo7744 Last updated by author(s): Jun 8, 2019 Reporting Summary Nature Research wishes to improve the reproducibility of the work that we publi...	4.6%
 2018-09-12741C-reporting summary.pdf https://s3.us-west-2.amazonaws.com/caltechauthors/81/13/465f-6af4-4dc6-a751-4fle0c7bc532/data?response... s.gribbin Reporting Summary Nature Research wishes to improve the reproducibility of the work that we publish. This form provides structure for con...	4.1%
 179913_2_related_ms_3465867_pk0499.pdf https://digitalcommons.unl.edu/cgi/viewcontent.cgi?filename=2&article=1375&context=virologypub&type=addit... mrp0730 Reporting Summary Nature Research wishes to improve the reproducibility of the work that we publish. This form provides structure for con...	4.1%
 Garcia_RS.pdf https://stacks.cdc.gov/view/cdc/80491/cdc_80491_ds17.pdf j.pamies Reporting Summary Nature Research wishes to improve the reproducibility of the work that we publish. This form provides structure for con...	4%
 208984_2_related_ms_4073340_pxwd5w.pdf https://pure.mpg.de/rest/items/item_3174783_3/component/file_3174792/content asp2065 Last updated by author(s): Sep 11, 2019 Reporting Summary Nature Research wishes to improve the reproducibility of the work that we publi...	3.7%
 3772_2_supp_109400_q9948d.pdf https://s3.us-west-2.amazonaws.com/caltechauthors/c5/41/7b0e-d234-4001-8e79-7e04ee6fld0c/data?respon... loq5070 Last updated by author(s): Apr 19, 2020 Reporting Summary Nature Research wishes to improve the reproducibility of the work that we publi...	3.6%

 198703_4_related_ms_4482495_q74k7v.pdf https://pure.mpg.de/rest/items/item_3220919/component/file_3220923/content mrp0730 Last updated by author(s): 9-3-2020 Reporting Summary Nature Research wishes to improve the reproducibility of the work that we publish....	3.6%
 Highly accurate protein structure prediction for the human proteome https://dasher.wustl.edu/bio5357/readings/nature-596-590-21.pdf Kathryn Tunyasuvunakool Article Highly accurate protein structure prediction for the human proteome https://doi.org/10.1038/s41586-021-03828-1 Received: 11 May ...	3%
 Highly Accurate Protein Structure Prediction with AlphaFold: Training with La... https://webeditions.page/works/highly-accurate-protein-structure-prediction-with-alphafold/06-training-data/ Skip to textTraining with labelled and unlabelled data ¶The AlphaFold architecture is able to train to high accuracy using only supervise...	2.8%
 (PDF) Accurate prediction of protein structures and interactions using a thre... https://www.researchgate.net/publication/353282026_accurate_prediction_of_protein_structures_and_interacti... ArticlePDF AvailableAccurate prediction of protein structures and interactions using a three-track neural network July 2021 Science ...	2.6%
 EconPapers: Highly accurate protein structure prediction with AlphaFold https://econpapers.repec.org/article/natnature/v_3a596_3ay_3a2021_3ai_3a7873_3ad_3a10.1038_5fs41586-021-... EconPapers Home About EconPape...	2.4%
 Highly accurate protein structure prediction with AlphaFold - Innovita Research https://www.innovitaresearch.com/2021/07/22/highly-accurate-protein-structure-prediction-with-alphafold/ Highly accurate protein structure prediction with AlphaFold Home News Related Science News General Highly accurate protein structure ...	2.2%
 OpenFold: retraining AlphaFold2 yields new insights into its learning mechani... https://www.nature.com/articles/s41592-024-02272-z Skip to main content Thank you for visiting nature.com. You are using a browser version with limited supp...	2.2%
 Fast, sensitive detection of protein homologs using deep dense retrieval Na... https://www.nature.com/articles/s41587-024-02353-6 Skip to main content Thank you for visiting nature.com. You are using a browser version with limited supp...	2.2%

-
-  Highly accurate protein structure prediction with AlphaFold - Seoul National ... 2.1%
<https://snu.elsevierpure.com/en/publications/highly-accurate-protein-structure-prediction-with-alpha-fold/>
Skip to main navigation Skip to search Skip to main content ...
-
-  Highly accurate protein structure prediction with AlphaFold - Seoul National ... 2.1%
<https://snu.elsevierpure.com/en/publications/highly-accurate-protein-structure-prediction-with-alpha-fold/>
Skip to main navigation Skip to search Skip to main content ...
-
-  RNA contact prediction by data efficient deep learning 2%
<https://publikationen.bibliothek.kit.edu/1000162205/151364829>
Oskar Taubert
ARTICLE <https://doi.org/10.1038/s42003-023-05244-9> OPEN RNA contact prediction by data efficient deep learning 1234567890();; Oskar ...
-
-  Highly accurate protein structure prediction with AlphaFold 2%
https://ideas.repec.org/a/nat/nature/v596y2021i7873d10.1038_s41586-021-03819-2.html
Browse Econ Literature Working pa...
-
-  heidelberg.ai 2%
<https://heidelberg.ai/2022/05/05/alpha-fold.html>
We are excited to have Simon Kohl, an alumni of the Medical Imaging Division of the DKFZ and co-found...
-
-  Jumper_2021_Nature_596_583 2%
https://bioweb.supagro.inrae.fr/esther/paper/abstract/jumper_2021_nature_596_583
"Proteins are essential to life, and understanding their structure can facilitate a mechanistic understanding of their function. Through ...
-
-  The Proteome Folding Problem and Cellular Proteostasis | Request PDF 2%
https://www.researchgate.net/publication/353884988_the_proteome_folding_problem_and_cellular_proteostasis
Home Molecular Biology Proteostasis ArticleLiterature ReviewThe Proteome Folding Problem and Cellular Proteostasis August 2021 Jour...
-
-  Being Brief. | Oxford Protein Informatics Group 2%
<https://www.blopig.com/blog/2021/09/being-brief/>
Skip to content This is a blog post about using fewer words. Start with incomplete sentence...

 ColabFold: making protein folding accessible to all https://storage.prod.researchhub.com/uploads/papers/2023/10/11/s41592-022-01488-1.pdf Milot Mirdita BRIEF COMMUNICATION https://doi.org/10.1038/s41592-022-01488-1 ColabFold: making protein folding accessible to all Milot Mirdita* *1,10 ...	1.8%
 Landgrebe_AlphaFold_ICBO2022.pdf https://icbo-conference.github.io/icbo2022/slides/landgrebe_alphafold_icbo2022.pdf What AlphaFold teaches us about deep learning using prior knowledge and ontologies Jobst Landgrebe September 27th ICBO 2022 Ann Arbor Con...	1.7%
 Predicting multiple conformations via sequence clustering and AlphaFold2 Na... https://www.nature.com/articles/s41586-023-06832-9 Skip to main content Thank you for visiting nature.com. You are using a browser version with limited supp...	1.7%
 Advances in protein structure prediction and design Nature Reviews Molecula... https://www.nature.com/articles/s41580-019-0163-x Skip to main content Thank you for visiting nature.com. You are using a browser version with limited supp...	1.5%
 AlphaFold predictions are valuable hypotheses and accelerate but do not repla... https://escholarship.org/content/qt58t2n89d/qt58t2n89d.pdf Thomas C. Terwilliger; Dorothee Liebschner; Tristan I. Croll; Christopher J. Williams; Airlie J. McCoy; Billy K. Poon; Pavel UC Berkeley UC Berkeley Previously Published Works Title AlphaFold predictions are valuable hypotheses and accelerate but do not replace ...	1.3%
 MassiveFold: unveiling AlphaFold's hidden potential with optimized and parall... https://www.nature.com/articles/s43588-024-00714-4 Skip to main content Thank you for visiting nature.com. You are using a browser version with limited supp...	1.2%
 Highly Accurate Protein Structure Prediction with AlphaFold: Related Work R... https://webeditions.page/works/highly-accurate-protein-structure-prediction-with-alphafold/09-related-work/ Skip to textRelated work ¶The prediction of protein structures has had a long and varied development, which is extensively covered in a n...	1.2%
 Array programming with NumPy Nature https://www.nature.com/articles/s41586-020-2649-2 Skip to main content Thank you for visiting nature.com. You are using a browser version with limited supp...	1.2%

 AlphaFold prediction of structural ensembles of disordered proteins | Nature ... 1.2%
<https://www.nature.com/articles/s41467-025-56572-9>
Skip to main content Thank you for visiting nature.com. You are using a browser version with limited supp...

 Structures of the Mycobacterium tuberculosis efflux pump EfpA reveal the mech... 1.1%
<https://www.osti.gov/pages/servlets/purl/2581825>
Shuhui Wang
Article <https://doi.org/10.1038/s41467-024-51948-9> Structures of the Mycobacterium tuberculosis efflux pump EfpA reveal the mechanisms ...

 A conditional protein diffusion model generates artificial programmable endon... 1.1%
<https://www.nature.com/articles/s41421-024-00728-2>
AI Source Match
Skip to main content Thank you for visiting nature.com. You are using a browser version with limited supp...

 File:AlphaFold 2.png - Wikimedia Commons 1.1%
https://commons.wikimedia.org/wiki/file:alphafold_2.png
Jump to content ...

 276792_3_related_ms_5738804_qw7777.pdf 1.1%
<https://archimer.ifremer.fr/doc/00755/86722/92167.pdf>
mrp0730
Please select the one below that is the best fit for your research. If you are not sure, read the appropriate sections before making your...

 Protein complex structure modeling by cross-modal alignment between cryo-EM m... 1.1%
<https://www.nature.com/articles/s41467-024-53116-5>
Skip to main content Thank you for visiting nature.com. You are using a browser version with limited supp...

 13476977.pdf 1%
<https://epub.jku.at/download/pdf/13476977.pdf>
AI Source Match
A DESIGN-BASED STUDY ON EMERGING TECHNOLOGIES IN MATHEMATICS TEACHER EDUCATION: THE CASE OF 3D MODELLING AND PRINTING (3DMP) Author Math...

 File:The performance of AlphaFold.png - Wikimedia Commons 1%
https://commons.wikimedia.org/wiki/file:the_performance_of_alphafold.png
Jump to content ...

 Architectural details a, Evoformer block. Arrows show the information... Do... https://www.researchgate.net/figure/architectural-details-a-evoformer-block-arrows-show-the-information-flo... Figure - available from: NatureThis content is subject to copyright. Terms and conditions apply. Download View publication Copy r...	1%
 ColabFold: making protein folding accessible to all Springer Nature Experim... https://experiments.springernature.com/articles/10.1038/s41592-022-01488-1 ...	1%
 Discovery, characterization, and comparative analysis of new UGT72 and UGT84 ... https://www.nature.com/articles/s42004-024-01231-1 AI Source Match Skip to main content Thank you for visiting nature.com. You are using a browser version with limited supp...	1%
 Molecular mechanism of siderophore regulation by the Pseudomonas aeruginosa B... https://www.nature.com/articles/s42003-024-05995-z AI Source Match Skip to main content Thank you for visiting nature.com. You are using a browser version with limited supp...	1%
 MOC1 cleaves Holliday junctions through a cooperative nick and counter-nick m... https://www.nature.com/articles/s41467-024-49490-9 AI Source Match Skip to main content Thank you for visiting nature.com. You are using a browser version with limited supp...	1%
 Unmasking AlphaFold to integrate experiments and predictions in multimeric co... https://www.nature.com/articles/s41467-024-52951-w AI Source Match Skip to main content Thank you for visiting nature.com. You are using a browser version with limited supp...	0.9%
 Early prediction of neoadjuvant chemotherapy response by exploiting a transfe... https://www.nature.com/articles/s41598-021-93592-z Skip to main content Thank you for visiting nature.com. You are using a browser version with limited supp...	0.9%
 A data-driven group retrosynthesis planning model inspired by neurosymbolic p... https://www.nature.com/articles/s41467-024-55374-9 AI Source Match Skip to main content Thank you for visiting nature.com. You are using a browser version with limited supp...	0.9%

 AlphaFold2 - SBGrid Documentation 0.8%
<https://docs.sbggrid.org/apps/alphafold2/>
Skip to content ...

 Automatic construction of Petri net models for computational simulations of m... 0.8%
<https://www.nature.com/articles/s41540-024-00464-z>
AI Source Match
Skip to main content Thank you for visiting nature.com. You are using a browser version with limited supp...

 Файл:Architectural details of AlphaFold 2.png – Википедия нэвтэрхий толь 0.8%
https://mn.wikipedia.org/wiki/%d0%a4%d0%b0%d0%b9%d0%bb:architectural_details_of_alphafold_2.png
Jump to content ...

 File:Architectural details of AlphaFold 2.png – Wikimedia Commons 0.8%
https://commons.wikimedia.org/wiki/file:architectural_details_of_alphafold_2.png
Jump to content ...

 Biological structure and function emerge from scaling unsupervised learning t... 0.8%
<https://www.biorxiv.org/content/10.1101/622803v3.full.pdf>
Alexander Rives, Joshua Meier, Tom Sercu, Siddharth Goyal, Zeming Lin, Demi Guo, Myle Ott, C. Lawrence Zitnick, Jerry
bioRxiv preprint doi: <https://doi.org/10.1101/622803>; this version posted August 31, 2020. The copyright holder for this preprint (which ...

 v1.pdf?c=1750403723000 0.7%
https://assets-eu.researchsquare.com/files/rs-6845168/v1_covered_56189380-b021-48a6-9f79-b0e4fc5015b8.p...
Boosting AlphaFold Protein Tertiary Structure Prediction through MSA Engineering and Extensive Model Sampling and Ranking in CASP16 Jianl...

 GSTU_Research_Compilation-2026.pdf 0.7%
https://www.gstu.edu.bd/dev/sidebar_items/uploaded_files/research_center/gstu_research_compilation-2026...
AI Source Match
i †MvcvjMÄ weÁvb I cÖhyw3 wek' we"vjq †MvcvjMÄ-8105, evsjv†k| ii mǝúv'bvq GwWUi-Bb-Pxd W. †gvt kvgmyj Av†iwdb, cwiPvjK, wimvP© †mǝUv...

 AlphaFold @ CASP13: "What just happened?" ‹ Some Thoughts on a Mysterious Uni... 0.6%
<https://moalquraishi.wordpress.com/2018/12/09/alphafold-casp13-what-just-happened/>
Update: An updated version of this blogpost was published as a (peer-reviewed) Letter to the Edito...

 State-of-the-RNArt: benchmarking current methods for RNA 3D structure predict...0.6%

<https://www.biorxiv.org/content/10.1101/2023.12.22.573067v3.full>

AI Source Match

Skip to main content ...

 Medium0.4%

<https://isamu-website.medium.com/what-are-the-most-exciting-areas-of-gnn-research-part-1-understanding-...>

SitemapOpen in app Sign in Medium Logo Write Search Sign in Isamu Isozaki I am a student. Alphafold G...

 ProHistoneDB | AOA2R6EFE60.4%

<https://www.prohistonedb.org/entry/a0a2r6efe6>

...

 A Systematic Review on the Generative AI Applications in Human Medical Genomics0.3%

<https://arxiv.org/pdf/2508.20275>

Anton Changalidis; Yury Barbitoff; Yulia Nasykhova; Andrey Glotov AI Source Match

A S YSTEMATIC R EVIEW ON THE G ENERATIVE AI A PPLICATIONS IN H UMAN M EDICAL G ENOMICS arXiv:2508.20275v1 [cs.LG] 27 Aug 2025 Anton Cha...

 References – AI4Science101 documentation0.3%

https://ai4science101.deepmodeling.com/en/latest/chapters/ai_for_scientific_discovery/references.html

Home Blog Tutorials Docs DeePMD-kit DP-GEN dpdata DPDispatcher Publications GitHub Home Blog Tutorials Docs Publications GitHub...

 References – AI4Science101 documentation0.3%

http://ai4science101.deepmodeling.com/en/latest/chapters/ai_for_scientific_discovery/references.html

AI Source Match

Home Blog Tutorials Docs DeePMD-kit DP-GEN dpdata DPDispatcher Publications GitHub Home Blog Tutorials Docs Publications GitHub...

 Michael Figurnov0.3%

<https://mfigurnov.github.io/>

AI Source Match

Michael Figurno...

 A Systematic Review on the Generative AI Applications in Human Medical Genomics0.3%

<https://arxiv.org/html/2508.20275v1>

AI Source Match

arXiv is now an independent nonprofit! Learn more License: CC BY 4.0 arXiv:2508.20275v1 [cs.L...

 20220208_Final_Thesis_corrections_approved_by_Internal.pdf	0.3%
https://pure.port.ac.uk/ws/portalfiles/portal/69121417/20220208_final_thesis_corrections_approved_by_interna...	
Rosie Graham	
Development of plastic degrading enzymes for future industrial PET biorecycling Rosie Graham The thesis is submitted in partial fulfilme...	
 Graph_Pattern_Matching_DEUTSCH_DOA27012022_AFV.pdf	0.3%
https://www.pure.ed.ac.uk/ws/portalfiles/portal/260606327/graph_pattern_matching_deutsch_doa27012022_a...	
Edinburgh Research Explorer Graph Pattern Matching in GQL and SQL/PGQ Citation for published version: Deutsch, A, Francis, N, Green, A, H...	
 High Accuracy Protein Structure Prediction Using Deep Learning	0.2%
https://sungsoo.github.io/2020/12/03/alphafold2.html	
Home Tags Categories Archive ...	
 PDB-Struct_A_Comprehensive_Benchmark_for_Structure-based_Protein_Design.pdf	0.2%
https://www.mlsb.io/papers_2023/pdb-struct_a_comprehensive_benchmark_for_structure-based_protein_desig...	
PDB-Struct: A Comprehensive Benchmark for Structure-based Protein Design Chuanrui Wang1,2 Bozitao Zhong 1,2 Zuobai Zhang1,2 Narendra Cha...	
 Sparse Annotation is Sufficient for Bootstrapping Dense Segmentation	0.2%
https://europepmc.org/api/getpdf?pmcid=pmc11195258	
Uri Manor	
bioRxiv preprint doi: https://doi.org/10.1101/2024.06.14.599135 ; this version posted October 26, 2024. The copyright holder for this prep...	
 af_backprop/alphafold/model/lddt.py · simondurr/ProteinMPNN at main	0.2%
https://huggingface.co/spaces/simondurr/proteinmpnn/blob/main/af_backprop/alphafold/model/lddt.py	
Fetching metadata from the HF Docker repository... Simon Duerr add fast af 85bd48b about...	
 IDDT Results	0.2%
https://swissmodel.expasy.org/lddt/deprecated/example_2/	
IDDT Local Distance Difference Test Examples Example 1 Example 2 Help Download This is the original implementation of IDDT...	
 References SWISS-MODEL	0.2%
https://swissmodel.expasy.org/docs/references	
References If you are using models from the SWISS-MODEL Server or Repository, please cite the corresponding articles:SWISS-MODEL Workspa...	

 content https://researchcommons.waikato.ac.nz/server/api/core/bitstreams/6e7afdbc-b5ee-41d0-93b2-9a3e4fe870ec... Hubbard, Amy http://researchcommons.waikato.ac.nz/ Research Commons at the University of Waikato Copyright Statement: The digital copy of this thesis ...	0.2%
 RCSB PDB - 6SJL: Structure of the Tle1 effector bound to the VgrG spike from ... https://www.rcsb.org/structure/6sjl 259,129 Structures from the PDB archive 1,062,058 Computed Structure Models (CSM) Structure Summary S...	0.2%
 Structural basis for loading and inhibition of a bacterial T6SS phospholipase... https://research.pasteur.fr/en/publication/structural-basis-for-loading-and-inhibition-of-a-bacterial-t6ss-phos... ...	0.2%
 XAI meets Biology: A Comprehensive Review of Explainable AI in Bioinformatics... https://arxiv.org/html/2312.06082v1 AI Source Match arXiv is now an independent nonprofit! Learn more License: CC BY-NC-ND 4.0 arXiv:2312.06082v1...	0.2%
 SciVideoBench: Benchmarking Scientific Video Reasoning in Large Multimodal Mo... https://arxiv.org/html/2510.08559v1 AI Source Match arXiv is now an independent nonprofit! Learn more License: CC BY 4.0 arXiv:2510.08559v1 [cs.C...	0.2%
 Triangle Multiplication is All You Need for Biomolecular Structure Representa... https://arxiv.org/html/2510.18870v1 arXiv is now an independent nonprofit! Learn more License: CC BY 4.0 arXiv:2510.18870v1 [q-bi...	0.2%
 d07379f3acf3af51dfc8598862cadfa0-Supplemental-Datasets_and_Benchmarks.pdf https://proceedings.neurips.cc/paper_files/paper/2023/file/d07379f3acf3af51dfc8598862cadfa0-supplemental... Symmetry-Informed Geometric Representation for Molecules, Proteins, and Crystalline Materials Shengchao Liu ^{1,2} , Weitao Du ³ , Yanjing Li...	0.2%
 Freddolino Lab - Chengxin Zhang https://freddolino-lab.med.umich.edu/people/chengxin-zhang Search this site Embedded Files Chengxin Zhang Research A...	0.1%

 ranganathan2005.pdf http://physics.bu.edu/~pankajm/py571/papers/ranganathan2005.pdf Vol 437 22 September 2005 doi:10.1038/nature03991 ARTICLES Evolutionary information for specifying a protein fold Michael Socolich1,2*, ...	0.1%
 Structure of SARS-CoV-2 ORF8, a rapidly evolving immune evasion protein PNAS https://www.pnas.org/doi/10.1073/pnas.2021785118 ...	0.1%
 IDDT: a local superposition-free score for comparing protein structures and m... https://edoc.unibas.ch/entities/publication/a380c334-f190-492b-9100-85b4650888b1 IDDT: a local superposition-free score for comparing protein structures and models using distance difference tests ...	0.1%
 Structure of SARS-CoV-2 ORF8, a rapidly evolving immune evasion protein (Jour... https://www.osti.gov/pages/biblio/1815732 ...	0.1%
 (IUCr) Protein structure prediction by AlphaFold2: are attention and symmetri... https://journals.iucr.org/d/issues/2021/08/00/rr5212/index.html AI Source Match Acta Crystallographica Section D Acta Crystallographica Section D STRUCTURAL BIOLOGY ...	0.1%
 The structure of human CST reveals a decameric assembly bound to telomeric DN... https://www.science.org/doi/10.1126/science.aaz9649 ...	0.1%
 Introduction – ProDy https://www.bahargroup.org/prody/tutorials/anmd_tutorial/intro.html index modules ...	0.1%
 Annotated Bibliography BIML https://berryvilleimi.com/bibliography/ AI Source Match ...	0.1%

Highly accurate protein structure prediction with AlphaFold


<https://doi.org/10.1038/s41586-021-03819-2>

Received: 11 May 2021

Accepted: 12 July 2021

Published online: 15 July 2021

Open access



John Jumper^{1,4}✉, Richard Evans^{1,4}, Alexander Pritzel^{1,4}, Tim Green^{1,4}, Michael Figurnov^{1,4}, Olaf Ronneberger^{1,4}, Kathryn Tunyasuvunakool^{1,4}, Russ Bates^{1,4}, Augustin Židek^{1,4}, Anna Potapenko^{1,4}, Alex Bridgland^{1,4}, Clemens Meyer^{1,4}, Simon A. A. Kohl^{1,4}, Andrew J. Ballard^{1,4}, Andrew Cowie^{1,4}, Bernardino Romera-Paredes^{1,4}, Stanislav Nikolov^{1,4}, Rishub Jain^{1,4}, Jonas Adler¹, Trevor Back¹, Stig Petersen¹, David Reiman¹, Ellen Clancy¹, Michal Zielinski, Martin Steinegger^{2,3}, Michalina Pacholska, Tamas Berghammer, Sebastian Bodenstein¹, David Silver¹, Oriol Vinyals¹, Andrew W. Senior¹, Koray Kavukcuoglu¹, Pushmeet Kohli¹ & Demis Hassabis^{1,4}✉

Proteins are essential to life, and understanding their structure can facilitate a mechanistic understanding of their function. Through an enormous experimental effort^{1–4}, the structures of around 100,000 unique proteins have been determined⁵, but this represents a small fraction of the billions of known protein sequences^{6,7}. Structural coverage is bottlenecked by the months to years of painstaking effort required to determine a single protein structure. Accurate computational approaches are needed to address this gap and to enable large-scale structural bioinformatics. Predicting the three-dimensional structure that a protein will adopt based solely on its amino acid sequence—the structure prediction component of the ‘protein folding problem’⁸—has been an important open research problem for more than 50 years⁹. Despite recent progress^{10–14}, existing methods fall far short of atomic accuracy, especially when no homologous structure is available. Here we provide the first computational method that can regularly predict protein structures with atomic accuracy even in cases in which no similar structure is known. We validated an entirely redesigned version of our neural network-based model, AlphaFold, in the challenging 14th Critical Assessment of protein Structure Prediction (CASP14)¹⁵, demonstrating accuracy competitive with experimental structures in a majority of cases and greatly outperforming other methods. Underpinning the latest version of AlphaFold is a novel machine learning approach that incorporates physical and biological knowledge about protein structure, leveraging multi-sequence alignments, into the design of the deep learning algorithm.

The development of computational methods to predict three-dimensional (3D) protein structures from the protein sequence has proceeded along two complementary paths that focus on either the physical interactions or the evolutionary history. The physical interaction programme heavily integrates our understanding of molecular driving forces into either thermodynamic or kinetic simulation of protein physics¹⁶ or statistical approximations thereof¹⁷. Although theoretically very appealing, this approach has proved highly challenging for even moderate-sized proteins due to the computational intractability of molecular simulation, the context dependence of protein stability and the difficulty of producing sufficiently accurate models of protein physics. The evolutionary programme has provided an alternative in recent years, in which the constraints on protein structure are derived from bioinformatics analysis of the evolutionary history of proteins, homology to solved structures^{18,19} and pairwise evolutionary correlations^{20–24}. This bioinformatics approach has benefited greatly from

the steady growth of experimental protein structures deposited in the Protein Data Bank (PDB)⁵, the explosion of genomic sequencing and the rapid development of deep learning techniques to interpret these correlations. Despite these advances, contemporary physical and evolutionary-history-based approaches produce predictions that are far short of experimental accuracy in the majority of cases in which a close homologue has not been solved experimentally and this has limited their utility for many biological applications.

In this study, we develop the first, to our knowledge, computational approach capable of predicting protein structures to near experimental accuracy in a majority of cases. The neural network AlphaFold that we developed was entered into the CASP14 assessment (May–July 2020; entered under the team name ‘AlphaFold2’ and a completely different model from our CASP13 AlphaFold system¹⁰). The CASP assessment is carried out biennially using recently solved structures that have not been deposited in the PDB or publicly disclosed so that it is a blind test

¹DeepMind, London, UK. ²School of Biological Sciences, Seoul National University, Seoul, South Korea. ³Artificial Intelligence Institute, Seoul National University, Seoul, South Korea. ⁴These authors contributed equally: John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Židek, Anna Potapenko, Alex Bridgland, Clemens Meyer, Simon A. A. Kohl, Andrew J. Ballard, Andrew Cowie, Bernardino Romera-Paredes, Stanislav Nikolov, Rishub Jain, Demis Hassabis.

✉e-mail: jumper@deepmind.com; dhcontact@deepmind.com




Fig. 1 | AlphaFold produces highly accurate structures. **a**, The performance of AlphaFold on the CASP14 dataset ($n = 87$ protein domains) relative to the top-15 entries (out of 146 entries), group numbers correspond to the numbers assigned to entrants by CASP. Data are median and the 95% confidence interval of the median, estimated from 10,000 bootstrap samples. **b**, Our prediction of CASP14 target T1049 (PDB 6Y4F, blue) compared with the true (experimental) structure (green). Four residues in the C terminus of the crystal structure are B-factor outliers and are not depicted. **c**, CASP14 target T1056 (PDB 6YJ1).

An example of a well-predicted zinc-binding site (AlphaFold has accurate side chains even though it does not explicitly predict the zinc ion). **d**, CASP target T1044 (PDB 6VR4)—a 2,180-residue single chain—was predicted with correct domain packing (the prediction was made after CASP using AlphaFold without intervention). **e**, Model architecture. Arrows show the information flow among the various components described in this paper. Array shapes are shown in parentheses with **s**, number of sequences (N_{seq} in the main text); **r**, number of residues (N_{res} in the main text); **c**, number of channels.

for the participating methods, and has long served as the gold-standard assessment for the accuracy of structure prediction^{25,26}.

In CASP14, AlphaFold structures were vastly more accurate than competing methods. AlphaFold structures had a median backbone accuracy of 0.96 Å r.m.s.d.₉₅ (C α root-mean-square deviation at 95% residue coverage) (95% confidence interval = 0.85–1.16 Å) whereas the next best performing method had a median backbone accuracy of 2.8 Å r.m.s.d.₉₅ (95% confidence interval = 2.7–4.0 Å) (measured on CASP domains; see Fig. 1a for backbone accuracy and Supplementary Fig. 14 for all-atom accuracy). As a comparison point for this accuracy, the width of a carbon atom is approximately 1.4 Å. In addition to very accurate domain structures (Fig. 1b), AlphaFold is able to produce highly accurate side chains (Fig. 1c) when the backbone is highly accurate and considerably improves over template-based methods even when strong templates are available. The all-atom accuracy of AlphaFold was 1.5 Å r.m.s.d.₉₅ (95% confidence interval = 1.2–1.6 Å) compared with the 3.5 Å r.m.s.d.₉₅ (95% confidence interval = 3.1–4.2 Å) of the best alternative method. Our methods are scalable to very long proteins with accurate domains and domain-packing (see Fig. 1d for the prediction of a 2,180-residue protein with no structural homologues). Finally, the model is able to provide precise, per-residue estimates of its reliability that should enable the confident use of these predictions.

We demonstrate in Fig. 2a that the high accuracy that AlphaFold demonstrated in CASP14 extends to a large sample of recently released PDB

structures; in this dataset, all structures were deposited in the PDB after our training data cut-off and are analysed as full chains (see Methods, Supplementary Fig. 15 and Supplementary Table 6 for more details). Furthermore, we observe high side-chain accuracy when the backbone prediction is accurate (Fig. 2b) and we show that our confidence measure, the predicted local-distance difference test (pLDDT), reliably predicts the $\text{C}\alpha$ local-distance difference test (IDDT- $\text{C}\alpha$) accuracy of the corresponding prediction (Fig. 2c). We also find that the global superposition metric template modelling score (TM-score)²⁷ can be accurately estimated (Fig. 2d). Overall, these analyses validate that the high accuracy and reliability of AlphaFold on CASP14 proteins also transfers to an uncured collection of recent PDB submissions, as would be expected (see Supplementary Methods 1.15 and Supplementary Fig. 11 for confirmation that this high accuracy extends to new folds).

The AlphaFold network

AlphaFold greatly improves the accuracy of structure prediction by incorporating novel neural network architectures and training procedures based on the evolutionary, physical and geometric constraints of protein structures. In particular, we demonstrate a new architecture to jointly embed multiple sequence alignments (MSAs) and pairwise features, a new output representation and associated loss that enable accurate end-to-end structure prediction, a new equivariant attention




Fig. 2 | Accuracy of AlphaFold on recent PDB structures. The analysed structures are newer than any structure in the training set. Further filtering is applied to reduce redundancy (see Methods). **a**, Histogram of backbone r.m.s.d. for full chains (Cα r.m.s.d. at 95% coverage). Error bars are 95% confidence intervals (Poisson). This dataset excludes proteins with a template (identified by hmmsearch) from the training set with more than 40% sequence identity covering more than 1% of the chain ($n = 3,144$ protein chains). The overall median is 1.46 Å (95% confidence interval = 1.40–1.56 Å). Note that this measure will be highly sensitive to domain packing and domain accuracy; a high r.m.s.d. is expected for some chains with uncertain packing or packing errors. **b**, Correlation between backbone accuracy and side-chain accuracy. Filtered to structures with any observed side chains and resolution better than 2.5 Å ($n = 5,317$ protein chains); side chains were further filtered to B-factor $< 30 \text{ Å}^2$. A rotamer is classified as correct if the predicted torsion angle is within 40° . Each point aggregates a range of IDDT-Cα, with a bin size of 2 units above 70 IDDT-Cα and 5 units otherwise. Points correspond to the mean accuracy; error bars are 95% confidence intervals (Student t-test) of the mean on a per-residue basis. **c**, Confidence score compared to the true accuracy on chains. Least-squares linear fit $\text{IDDT-C}\alpha = 0.997 \times \text{pLDDT} - 1.17$ (Pearson's $r = 0.76$). $n = 10,795$ protein chains. The shaded region of the linear fit represents a 95% confidence interval estimated from 10,000 bootstrap samples. In the companion paper³⁹, additional quantification of the reliability of pLDDT as a confidence measure is provided. **d**, Correlation between pTM and full chain TM-score. Least-squares linear fit $\text{TM-score} = 0.98 \times \text{pTM} + 0.07$ (Pearson's $r = 0.85$). $n = 10,795$ protein chains. The shaded region of the linear fit represents a 95% confidence interval estimated from 10,000 bootstrap samples.

architecture, use of intermediate losses to achieve iterative refinement of predictions, masked MSA loss to jointly train with the structure, learning from unlabelled protein sequences using self-distillation and self-estimates of accuracy.

The AlphaFold network directly predicts the 3D coordinates of all heavy atoms for a given protein using the primary amino acid sequence and aligned sequences of homologues as inputs (Fig. 1e; see Methods for details of inputs including databases, MSA construction and use of templates). A description of the most important ideas and components is provided below. The full network architecture and training procedure are provided in the Supplementary Methods.

The network comprises two main stages. First, the trunk of the network processes the inputs through repeated layers of a novel neural network block that we term Evoformer to produce an $N_{\text{seq}} \times N_{\text{res}}$ array (N_{seq} , number of sequences; N_{res} , number of residues) that represents a processed MSA and an $N_{\text{res}} \times N_{\text{res}}$ array that represents residue pairs. The MSA representation is initialized with the raw MSA (although see Supplementary Methods 1.2.7 for details of handling very deep MSAs). The Evoformer blocks contain a number of attention-based and non-attention-based components. We show evidence in ‘Interpreting the neural network’ that a concrete structural hypothesis arises early within the Evoformer blocks and is continuously refined. The key innovations in the Evoformer block are new mechanisms to exchange information within the MSA and pair representations that enable direct reasoning about the spatial and evolutionary relationships.

The trunk of the network is followed by the structure module that introduces an explicit 3D structure in the form of a rotation and translation for each residue of the protein (global rigid body frames). These representations are initialized in a trivial state with all rotations set to the identity and all positions set to the origin, but rapidly develop and refine a highly accurate protein structure with precise atomic details. Key innovations in this section of the network include breaking the chain structure to allow simultaneous local refinement of all parts of the structure, a novel equivariant transformer to allow the network to implicitly reason about the unrepresented side-chain atoms and a loss term that places substantial weight on the orientational correctness of the residues. Both within the structure module and throughout the whole network, we reinforce the notion of iterative refinement by repeatedly applying the final loss to outputs and then feeding the outputs recursively into the same modules. The iterative refinement using the whole network (which we term ‘recycling’ and is related to approaches in computer vision^{28,29}) contributes markedly to accuracy with minor extra training time (see Supplementary Methods 1.8 for details).

Evoformer

The key principle of the building block of the network—named Evoformer (Figs. 1e, 3a)—is to view the prediction of protein structures as a graph inference problem in 3D space in which the edges of the graph are defined by residues in proximity. The elements of the pair representation encode information about the relation between the residues (Fig. 3b). The columns of the MSA representation encode the individual residues of the input sequence while the rows represent the sequences in which those residues appear. Within this framework, we define a number of update operations that are applied in each block in which the different update operations are applied in series.

The MSA representation updates the pair representation through an element-wise outer product that is summed over the MSA sequence dimension. In contrast to previous work³⁰, this operation is applied within every block rather than once in the network, which enables the continuous communication from the evolving MSA representation to the pair representation.

Within the pair representation, there are two different update patterns. Both are inspired by the necessity of consistency of the pair




Fig. 3 | Architectural details. **a**, Evoformer block. Arrows show the information flow. The shape of the arrays is shown in parentheses. **b**, The pair representation interpreted as directed edges in a graph. **c**, Triangle multiplicative update and triangle self-attention. The circles represent residues. Entries in the pair representation are illustrated as directed edges and in each diagram, the edge being updated is ij . **d**, Structure module including Invariant point attention (IPA)

module. The single representation is a copy of the first row of the MSA representation. **e**, Residue gas: a representation of each residue as one free-floating rigid body for the backbone (blue triangles) and γ angles for the side chains (green circles). The corresponding atomic structure is shown below. **f**, Frame aligned point error (FAPE). Green, predicted structure; grey, true structure; (R_k, t_k) , frames; x_i , atom positions.

representation—for a pairwise description of amino acids to be representable as a single 3D structure, many constraints must be satisfied including the triangle inequality on distances. On the basis of this intuition, we arrange the update operations on the pair representation in terms of triangles of edges involving three different nodes (Fig. 3c). In particular, we add an extra logit bias to axial attention³¹ to include the ‘missing edge’ of the triangle and we define a non-attention update operation ‘triangle multiplicative update’ that uses two edges to update the missing third edge (see Supplementary Methods 1.6.5 for details). The triangle multiplicative update was developed originally as a more symmetric and cheaper replacement for the attention, and networks that use only the attention or multiplicative update are both able to produce high-accuracy structures. However, the combination of the two updates is more accurate.

We also use a variant of axial attention within the MSA representation. During the per-sequence attention in the MSA, we project additional logits from the pair stack to bias the MSA attention. This closes the loop by providing information flow from the pair representation back into the MSA representation, ensuring that the overall Evoformer block is able to fully mix information between the pair and MSA representations and prepare for structure generation within the structure module.

End-to-end structure prediction

The structure module (Fig. 3d) operates on a concrete 3D backbone structure using the pair representation and the original sequence row (single representation) of the MSA representation from the trunk. The 3D backbone structure is represented as N_{res} independent rotations and translations, each with respect to the global frame (residue gas) (Fig. 3e). These rotations and translations—representing the geometry of the N-Ca-C atoms—prioritize the orientation of the protein backbone so that the location of the side chain of each residue is highly constrained within that frame. Conversely, the peptide bond geometry is completely unconstrained and the network is observed to frequently violate the chain constraint during the application of the structure module as breaking this constraint enables the local refinement of all parts of the chain without solving complex loop closure problems. Satisfaction of the peptide bond geometry is encouraged during fine-tuning by a violation loss term. Exact enforcement of peptide bond geometry is only achieved in the post-prediction relaxation of the structure by gradient descent in the Amber³² force field. Empirically, this final relaxation does not improve the accuracy of the model as measured by the

global distance test (GDT)³³ or IDDT-C α ³⁴ but does remove distracting stereochemical violations without the loss of accuracy.

The residue gas representation is updated iteratively in two stages (Fig. 3d). First, a geometry-aware attention operation that we term ‘invariant point attention’ (IPA) is used to update an N_{res} set of neural activations (single representation) without changing the 3D positions, then an equivariant update operation is performed on the residue gas using the updated activations. The IPA augments each of the usual attention queries, keys and values with 3D points that are produced in the local frame of each residue such that the final value is invariant to global rotations and translations (see Methods ‘IPA’ for details). The 3D queries and keys also impose a strong spatial/locality bias on the attention, which is well-suited to the iterative refinement of the protein structure. After each attention operation and element-wise transition block, the module computes an update to the rotation and translation of each backbone frame. The application of these updates within the local frame of each residue makes the overall attention and update block an equivariant operation on the residue gas.

Predictions of side-chain χ angles as well as the final, per-residue accuracy of the structure (pLDDT) are computed with small per-residue networks on the final activations at the end of the network. The estimate of the TM-score (pTM) is obtained from a pairwise error prediction that is computed as a linear projection from the final pair representation. The final loss (which we term the frame-aligned point error (FAPE) (Fig. 3f)) compares the predicted atom positions to the true positions under many different alignments. For each alignment, defined by aligning the predicted frame (R_k, t_k) to the corresponding true frame, we compute the distance of all predicted atom positions $\mathbf{x}_i$ from the true atom positions. The resulting $N_{\text{frames}} \times N_{\text{atoms}}$ distances are penalized with a clamped L_1 loss. This creates a strong bias for atoms to be correct relative to the local frame of each residue and hence correct with respect to its side-chain interactions, as well as providing the main source of chirality for AlphaFold (Supplementary Methods 1.9.3 and Supplementary Fig. 9).

Training with labelled and unlabelled data

The AlphaFold architecture is able to train to high accuracy using only supervised learning on PDB data, but we are able to enhance accuracy (Fig. 4a) using an approach similar to noisy student self-distillation³⁵. In this procedure, we use a trained network to predict the structure of around 350,000 diverse sequences from UniClust30³⁶ and make a new dataset of predicted structures filtered to a high-confidence subset. We then train the same architecture again from scratch using a mixture of PDB data and this new dataset of predicted structures as the training data, in which the various training data augmentations such as cropping and MSA subsampling make it challenging for the network to recapitulate the previously predicted structures. This self-distillation procedure makes effective use of the unlabelled sequence data and considerably improves the accuracy of the resulting network.

Additionally, we randomly mask out or mutate individual residues within the MSA and have a Bidirectional Encoder Representations from Transformers (BERT)-style³⁷ objective to predict the masked elements of the MSA sequences. This objective encourages the network to learn to interpret phylogenetic and covariation relationships without hardcoding a particular correlation statistic into the features. The BERT objective is trained jointly with the normal PDB structure loss on the same training examples and is not pre-trained, in contrast to recent independent work³⁸.

Interpreting the neural network

To understand how AlphaFold predicts protein structure, we trained a separate structure module for each of the 48 Evoformer blocks in the network while keeping all parameters of the main network frozen (Supplementary Methods 1.14). Including our recycling stages, this provides a trajectory of 192 intermediate structures—one per full




Fig. 4 | Interpreting the neural network. **a**, Ablation results on two target sets: the CASP14 set of domains ($n = 87$ protein domains) and the PDB test set of chains with template coverage of $\leq 30\%$ at 30% identity ($n = 2,261$ protein chains). Domains are scored with GDT and chains are scored with IDDT-C α . The ablations are reported as a difference compared with the average of the three baseline seeds. Means (points) and 95% bootstrap percentile intervals (error bars) are computed using bootstrap estimates of 10,000 samples. **b**, Domain GDT trajectory over 4 recycling iterations and 48 Evoformer blocks on CASP14 targets LmrP (T1024) and Orf8 (T1064) where D1 and D2 refer to the individual domains as defined by the CASP assessment. Both T1024 domains obtain the correct structure early in the network, whereas the structure of T1064 changes multiple times and requires nearly the full depth of the network to reach the final structure. Note, 48 Evoformer blocks comprise one recycling iteration.

Evoformer block—in which each intermediate represents the belief of the network of the most likely structure at that block. The resulting trajectories are surprisingly smooth after the first few blocks, showing that AlphaFold makes constant incremental improvements to the structure until it can no longer improve (see Fig. 4b for a trajectory of accuracy). These trajectories also illustrate the role of network depth. For very challenging proteins such as ORF8 of SARS-CoV-2 (T1064), the network searches and rearranges secondary structure elements for many layers before settling on a good structure. For other proteins such as LmrP (T1024), the network finds the final structure within the first few layers. Structure trajectories of CASP14 targets T1024, T1044, T1064 and T1091 that demonstrate a clear iterative building process for a range of protein sizes and difficulties are shown in Supplementary Videos 1–4. In Supplementary Methods 1.16 and Supplementary Figs. 12, 13, we interpret the attention maps produced by AlphaFold layers.

Figure 4a contains detailed ablations of the components of AlphaFold that demonstrate that a variety of different mechanisms contribute to AlphaFold accuracy. Detailed descriptions of each ablation model, their training details, extended discussion of ablation results and the




Fig. 5 | Effect of MSA depth and cross-chain contacts. **a**, Backbone accuracy (IDDT-C α) for the redundancy-reduced set of the PDB after our training data cut-off, restricting to proteins in which at most 25% of the long-range contacts are between different heteromer chains. We further consider two groups of proteins based on template coverage at 30% sequence identity: covering more than 60% of the chain ($n = 6,743$ protein chains) and covering less than 30% of the chain ($n = 1,596$ protein chains). MSA depth is computed by counting the

number of non-gap residues for each position in the MSA (using the N_{eff} weighting scheme; see Methods for details) and taking the median across residues. The curves are obtained through Gaussian kernel average smoothing (window size is 0.2 units in $\log_{10}(N_{\text{eff}})$); the shaded area is the 95% confidence interval estimated using bootstrap of 10,000 samples. **b**, An intertwined homotrimer (PDB 6SK0) is correctly predicted without input stoichiometry and only a weak template (blue is predicted and green is experimental).

effect of MSA depth on each ablation are provided in Supplementary Methods 1.13 and Supplementary Fig. 10.

MSA depth and cross-chain contacts

Although AlphaFold has a high accuracy across the vast majority of deposited PDB structures, we note that there are still factors that affect accuracy or limit the applicability of the model. The model uses MSAs and the accuracy decreases substantially when the median alignment depth is less than around 30 sequences (see Fig. 5a for details). We observe a threshold effect where improvements in MSA depth over around 100 sequences lead to small gains. We hypothesize that the MSA information is needed to coarsely find the correct structure within the early stages of the network, but refinement of that prediction into a high-accuracy model does not depend crucially on the MSA information. The other substantial limitation that we have observed is that AlphaFold is much weaker for proteins that have few intra-chain or homotypic contacts compared to the number of heterotypic contacts (further details are provided in a companion paper³⁹). This typically occurs for bridging domains within larger complexes in which the shape of the protein is created almost entirely by interactions with other chains in the complex. Conversely, AlphaFold is often able to give high-accuracy predictions for homomers, even when the chains are substantially intertwined (Fig. 5b). We expect that the ideas of AlphaFold are readily applicable to predicting full hetero-complexes in a future system and that this will remove the difficulty with protein chains that have a large number of hetero-contacts.

Related work

The prediction of protein structures has had a long and varied development, which is extensively covered in a number of reviews^{14,40–43}. Despite the long history of applying neural networks to structure prediction^{14,42,43}, they have only recently come to improve structure prediction^{10,11,44,45}. These approaches effectively leverage the rapid improvement in computer vision systems⁴⁶ by treating the problem of protein structure prediction as converting an ‘image’ of evolutionary couplings^{22–24} to an ‘image’ of the protein distance matrix and then integrating the distance predictions into a heuristic system that produces the final 3D coordinate prediction. A few recent studies have been developed to predict the 3D coordinates directly^{47–50}, but the accuracy of these approaches does not

match traditional, hand-crafted structure prediction pipelines⁵¹. In parallel, the success of attention-based networks for language processing⁵² and, more recently, computer vision^{31,53} has inspired the exploration of attention-based methods for interpreting protein sequences^{54–56}.

Discussion

The methodology that we have taken in designing AlphaFold is a combination of the bioinformatics and physical approaches: we use a physical and geometric inductive bias to build components that learn from PDB data with minimal imposition of handcrafted features (for example, AlphaFold builds hydrogen bonds effectively without a hydrogen bond score function). This results in a network that learns far more efficiently from the limited data in the PDB but is able to cope with the complexity and variety of structural data.

In particular, AlphaFold is able to handle missing the physical context and produce accurate models in challenging cases such as intertwined homomers or proteins that only fold in the presence of an unknown haem group. The ability to handle underspecified structural conditions is essential to learning from PDB structures as the PDB represents the full range of conditions in which structures have been solved. In general, AlphaFold is trained to produce the protein structure most likely to appear as part of a PDB structure. For example, in cases in which a particular stoichiometry, ligand or ion is predictable from the sequence alone, AlphaFold is likely to produce a structure that respects those constraints implicitly.

AlphaFold has already demonstrated its utility to the experimental community, both for molecular replacement⁵⁷ and for interpreting cryogenic electron microscopy maps⁵⁸. Moreover, because AlphaFold outputs protein coordinates directly, AlphaFold produces predictions in graphics processing unit (GPU) minutes to GPU hours depending on the length of the protein sequence (for example, around one GPU minute per model for 384 residues; see Methods for details). This opens up the exciting possibility of predicting structures at the proteome-scale and beyond—in a companion paper³⁹, we demonstrate the application of AlphaFold to the entire human proteome³⁹.

The explosion in available genomic sequencing techniques and data has revolutionized bioinformatics but the intrinsic challenge of experimental structure determination has prevented a similar expansion in our structural knowledge. By developing an accurate protein structure

prediction algorithm, coupled with existing large and well-curated structure and sequence databases assembled by the experimental community, we hope to accelerate the advancement of structural bioinformatics that can keep pace with the genomics revolution. We hope that AlphaFold—and computational approaches that apply its techniques for other biophysical problems—will become essential tools of modern biology.

Online content

Any methods, additional references, Nature Research reporting summaries, source data, extended data, supplementary information, acknowledgements, peer review information; details of author contributions and competing interests; and statements of data and code availability are available at <https://doi.org/10.1038/s41586-021-03819-2>.

1. Thompson, M. C., Yeates, T. O. & Rodriguez, J. A. Advances in methods for atomic resolution macromolecular structure determination. *F1000Res* **9**, 667 (2020).
2. Bai, X.-C., McMullan, G. & Scheres, S. H. W. How cryo-EM is revolutionizing structural biology. *Trends Biochem. Sci.* **40**, 49–57 (2015).
3. Jaskolski, M., Dauter, Z. & Wlodawer, A. A brief history of macromolecular crystallography. Illustrated by a family tree and its Nobel fruits. *FEBS J.* **281**, 3985–4009 (2014).
4. Wüthrich, K. The way to NMR structures of proteins. *Nat. Struct. Biol.* **8**, 923–925 (2001).
5. wwPDB Consortium. Protein Data Bank: the single global archive for 3D macromolecular structure data. *Nucleic Acids Res.* **47**, D520–D528 (2018).
6. Mitchell, A. L. et al. MGnify: the microbiome analysis resource in 2020. *Nucleic Acids Res.* **48**, D570–D578 (2020).
7. Steingger, M., Mirdita, M. & Söding, J. Protein-level assembly increases protein sequence recovery from metagenomic samples manyfold. *Nat. Methods* **16**, 603–606 (2019).
8. Dill, K. A., Ozkan, S. B., Shell, M. S. & Weikl, T. R. The protein folding problem. *Annu. Rev. Biophys.* **37**, 289–316 (2008).
9. Anfinsen, C. B. Principles that govern the folding of protein chains. *Science* **181**, 223–230 (1973).
10. Senior, A. W. et al. Improved protein structure prediction using potentials from deep learning. *Nature* **577**, 706–710 (2020).
11. Wang, S., Sun, S., Li, Z., Zhang, R. & Xu, J. Accurate de novo prediction of protein contact map by ultra-deep learning model. *PLOS Comput. Biol.* **13**, e1005324 (2017).
12. Zheng, W. et al. Deep-learning contact-map guided protein structure prediction in CASP13. *Proteins* **87**, 1149–1164 (2019).
13. Abriata, L. A., Tamó, G. E. & Dal Peraro, M. A further leap of improvement in tertiary structure prediction in CASP13 prompts new routes for future assessments. *Proteins* **87**, 1100–1112 (2019).
14. Pearce, R. & Zhang, Y. Deep learning techniques have significantly impacted protein structure prediction and protein design. *Curr. Opin. Struct. Biol.* **68**, 194–207 (2021).
15. Moul, J., Fidelis, K., Kryshtafovych, A., Schwede, T. & Topf, M. Critical assessment of techniques for protein structure prediction, fourteenth round. *CASP 14 Abstract Book* https://www.predictioncenter.org/casp14/doc/CASP14_Abstracts.pdf (2020).
16. Brini, E., Simmerling, C. & Dill, K. Protein storytelling through physics. *Science* **370**, eaaz3041 (2020).
17. Sippl, M. J. Calculation of conformational ensembles from potentials of mean force. An approach to the knowledge-based prediction of local structures in globular proteins. *J. Mol. Biol.* **213**, 859–883 (1990).
18. Sali, A. & Blundell, T. L. Comparative protein modelling by satisfaction of spatial restraints. *J. Mol. Biol.* **234**, 779–815 (1993).
19. Roy, A., Kucukural, A. & Zhang, Y. I-TASSER: a unified platform for automated protein structure and function prediction. *Nat. Protocols* **5**, 725–738 (2010).
20. Altschuh, D., Lesk, A. M., Bloomer, A. C. & Klug, A. Correlation of co-ordinated amino acid substitutions with function in viruses related to tobacco mosaic virus. *J. Mol. Biol.* **193**, 693–707 (1987).
21. Shindyalov, I. N., Kolchanov, N. A. & Sander, C. Can three-dimensional contacts in protein structures be predicted by analysis of correlated mutations? *Protein Eng.* **7**, 349–358 (1994).
22. Weigt, M., White, R. A., Szurmant, H., Hoch, J. A. & Hwa, T. Identification of direct residue contacts in protein–protein interaction by message passing. *Proc. Natl Acad. Sci. USA* **106**, 67–72 (2009).
23. Marks, D. S. et al. Protein 3D structure computed from evolutionary sequence variation. *PLoS ONE* **6**, e28766 (2011).
24. Jones, D. T., Buchan, D. W. A., Cozzetto, D. & Pontil, M. PSICOV: precise structural contact prediction using sparse inverse covariance estimation on large multiple sequence alignments. *Bioinformatics* **28**, 184–190 (2012).
25. Moul, J., Pedersen, J. T., Judson, R. & Fidelis, K. A large-scale experiment to assess protein structure prediction methods. *Proteins* **23**, ii–iv (1995).
26. Kryshtafovych, A., Schwede, T., Topf, M., Fidelis, K. & Moul, J. Critical assessment of methods of protein structure prediction (CASP)-round XIII. *Proteins* **87**, 1011–1020 (2019).
27. Zhang, Y. & Skolnick, J. Scoring function for automated assessment of protein structure template quality. *Proteins* **57**, 702–710 (2004).
28. Tu, Z. & Bai, X. Auto-context and its application to high-level vision tasks and 3D brain image segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.* **32**, 1744–1757 (2010).
29. Carreira, J., Agrawal, P., Fragkiadaki, K. & Malik, J. Human pose estimation with iterative error feedback. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition* 4733–4742 (2016).
30. Mirabello, C. & Wallner, B. rawMSA: end-to-end deep learning using raw multiple sequence alignments. *PLoS ONE* **14**, e0220182 (2019).
31. Huang, Z. et al. CCNet: criss-cross attention for semantic segmentation. In *Proc. IEEE/CVF International Conference on Computer Vision* 603–612 (2019).
32. Hornak, V. et al. Comparison of multiple Amber force fields and development of improved protein backbone parameters. *Proteins* **65**, 712–725 (2006).
33. Zemla, A. LGA: a method for finding 3D similarities in protein structures. *Nucleic Acids Res.* **31**, 3370–3374 (2003).
34. Mariani, V., Biasini, M., Barbato, A. & Schwede, T. IDDT: a local superposition-free score for comparing protein structures and models using distance difference tests. *Bioinformatics* **29**, 2722–2728 (2013).
35. Xie, Q., Luong, M.-T., Hovy, E. & Le, Q. V. Self-training with noisy student improves imagenet classification. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition* 10687–10698 (2020).
36. Mirdita, M. et al. Uniclust databases of clustered and deeply annotated protein sequences and alignments. *Nucleic Acids Res.* **45**, D170–D176 (2017).
37. Devlin, J., Chang, M.-W., Lee, K. & Toutanova, K. BERT: pre-training of deep bidirectional transformers for language understanding. In *Proc. 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies* 1, 4171–4186 (2019).
38. Rao, R. et al. MSA transformer. In *Proc. 38th International Conference on Machine Learning* 139, 8844–8856 (2021).
39. Tunyasuvunakool, K. et al. Highly accurate protein structure prediction for the human proteome. *Nature* <https://doi.org/10.1038/s41586-021-03828-1> (2021).
40. Kuhlman, B. & Bradley, P. Advances in protein structure prediction and design. *Nat. Rev. Mol. Cell Biol.* **20**, 681–697 (2019).
41. Marks, D. S., Hopf, T. A. & Sander, C. Protein structure prediction from sequence variation. *Nat. Biotechnol.* **30**, 1072–1080 (2012).
42. Qian, N. & Sejnowski, T. J. Predicting the secondary structure of globular proteins using neural network models. *J. Mol. Biol.* **202**, 865–884 (1988).
43. Fariselli, P., Olmea, O., Valencia, A. & Casadio, R. Prediction of contact maps with neural networks and correlated mutations. *Protein Eng.* **14**, 835–843 (2001).
44. Yang, J. et al. Improved protein structure prediction using predicted interresidue orientations. *Proc. Natl Acad. Sci. USA* **117**, 1496–1503 (2020).
45. Li, Y. et al. Deducing high-accuracy protein contact-maps from a triplet of coevolutionary matrices through deep residual convolutional networks. *PLOS Comput. Biol.* **17**, e1008865 (2021).
46. He, K., Zhang, X., Ren, S. & Sun, J. Deep residual learning for image recognition. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition* 770–778 (2016).
47. AlQuraishi, M. End-to-end differentiable learning of protein structure. *Cell Syst.* **8**, 292–301 (2019).
48. Senior, A. W. et al. Protein structure prediction using multiple deep neural networks in the 13th Critical Assessment of Protein Structure Prediction (CASP13). *Proteins* **87**, 1141–1148 (2019).
49. Ingraham, J., Riesselman, A. J., Sander, C. & Marks, D. S. Learning protein structure with a differentiable simulator. In *Proc. International Conference on Learning Representations* (2019).
50. Li, J. Universal transforming geometric network. Preprint at <https://arxiv.org/abs/1908.00723> (2019).
51. Xu, J., McPartlon, M. & Li, J. Improved protein structure prediction by deep learning irrespective of co-evolution information. *Nat. Mach. Intell.* **3**, 601–609 (2021).
52. Vaswani, A. et al. Attention is all you need. In *Advances in Neural Information Processing Systems* 5998–6008 (2017).
53. Wang, H. et al. Axial-deeplab: stand-alone axial-attention for panoptic segmentation. In *European Conference on Computer Vision* 108–126 (Springer, 2020).
54. Alley, E. C., Khimulya, G., Biswas, S., AlQuraishi, M. & Church, G. M. Unified rational protein engineering with sequence-based deep representation learning. *Nat. Methods* **16**, 1315–1322 (2019).
55. Heinzinger, M. et al. Modeling aspects of the language of life through transfer-learning protein sequences. *BMC Bioinformatics* **20**, 723 (2019).
56. Rives, A. et al. Biological structure and function emerge from scaling unsupervised learning to 250 million protein sequences. *Proc. Natl Acad. Sci. USA* **118**, e2016239118 (2021).
57. Pereira, J. et al. High-accuracy protein structure prediction in CASP14. *Proteins* <https://doi.org/10.1002/prot.26171> (2021).
58. Gupta, M. et al. CryoEM and AI reveal a structure of SARS-CoV-2 Nsp2, a multifunctional protein involved in key host processes. Preprint at <https://doi.org/10.1101/2021.05.10.443524> (2021).

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2021

Methods

Full algorithm details

Extensive explanations of the components and their motivations are available in Supplementary Methods 1.1–1.10, in addition, pseudocode is available in Supplementary Information Algorithms 1–32, network diagrams in Supplementary Figs. 1–8, input features in Supplementary Table 1 and additional details are provided in Supplementary Tables 2, 3. Training and inference details are provided in Supplementary Methods 1.11–1.12 and Supplementary Tables 4, 5.

IPA

The IPA module combines the pair representation, the single representation and the geometric representation to update the single representation (Supplementary Fig. 8). Each of these representations contributes affinities to the shared attention weights and then uses these weights to map its values to the output. The IPA operates in 3D space. Each residue produces query points, key points and value points in its local frame. These points are projected into the global frame using the backbone frame of the residue in which they interact with each other. The resulting points are then projected back into the local frame. The affinity computation in the 3D space uses squared distances and the coordinate transformations ensure the invariance of this module with respect to the global frame (see Supplementary Methods 1.8.2 ‘Invariant point attention (IPA)’ for the algorithm, proof of invariance and a description of the full multi-head version). A related construction that uses classic geometric invariants to construct pairwise features in place of the learned 3D points has been applied to protein design⁵⁹.

In addition to the IPA, standard dot product attention is computed on the abstract single representation and a special attention on the pair representation. The pair representation augments both the logits and the values of the attention process, which is the primary way in which the pair representation controls the structure generation.

Inputs and data sources

Inputs to the network are the primary sequence, sequences from evolutionarily related proteins in the form of a MSA created by standard tools including jackhmmer⁶⁰ and HHblits⁶¹, and 3D atom coordinates of a small number of homologous structures (templates) where available. For both the MSA and templates, the search processes are tuned for high recall; spurious matches will probably appear in the raw MSA but this matches the training condition of the network.

One of the sequence databases used, Big Fantastic Database (BFD), was custom-made and released publicly (see ‘Data availability’) and was used by several CASP teams. BFD is one of the largest publicly available collections of protein families. It consists of 65,983,866 families represented as MSAs and hidden Markov models (HMMs) covering 2,204,359,010 protein sequences from reference databases, metagenomes and metatranscriptomes.

BFD was built in three steps. First, 2,423,213,294 protein sequences were collected from UniProt (Swiss-Prot&TrEMBL, 2017-11)⁶², a soil reference protein catalogue and the marine eukaryotic reference catalogue⁷, and clustered to 30% sequence identity, while enforcing a 90% alignment coverage of the shorter sequences using MMseqs2/Linclust⁶³. This resulted in 345,159,030 clusters. For computational efficiency, we removed all clusters with less than three members, resulting in 61,083,719 clusters. Second, we added 166,510,624 representative protein sequences from Metaclust NR (2017-05; discarding all sequences shorter than 150 residues)⁶³ by aligning them against the cluster representatives using MMseqs2⁶⁴. Sequences that fulfilled the sequence identity and coverage criteria were assigned to the best scoring cluster. The remaining 25,347,429 sequences that could not be assigned were clustered separately and added as new clusters, resulting in the final clustering. Third, for each of the clusters, we computed an MSA using

FAMSA⁶⁵ and computed the HMMs following the Uniclust HH-suite database protocol³⁶.

The following versions of public datasets were used in this study. Our models were trained on a copy of the PDB⁵ downloaded on 28 August 2019. For finding template structures at prediction time, we used a copy of the PDB downloaded on 14 May 2020, and the PDB70⁶⁶ clustering database downloaded on 13 May 2020. For MSA lookup at both training and prediction time, we used Uniref90⁶⁷ v.2020_01, BFD, Uniclust30³⁶ v.2018_08 and MGnify⁶ v.2018_12. For sequence distillation, we used Uniclust30³⁶ v.2018_08 to construct a distillation structure dataset. Full details are provided in Supplementary Methods 1.2.

For MSA search on BFD + Uniclust30, and template search against PDB70, we used HHblits⁶¹ and HHSearch⁶⁶ from hh-suite v.3.0-beta.3 (version 14/07/2017). For MSA search on Uniref90 and clustered MGnify, we used jackhmmer from HMMER3⁶⁸. For constrained relaxation of structures, we used OpenMM v.7.3.1⁶⁹ with the Amber99sb force field³². For neural network construction, running and other analyses, we used TensorFlow⁷⁰, Sonnet⁷¹, NumPy⁷², Python⁷³ and Colab⁷⁴.

To quantify the effect of the different sequence data sources, we re-ran the CASP14 proteins using the same models but varying how the MSA was constructed. Removing BFD reduced the mean accuracy by 0.4 GDT, removing MGnify reduced the mean accuracy by 0.7 GDT, and removing both reduced the mean accuracy by 6.1 GDT. In each case, we found that most targets had very small changes in accuracy but a few outliers had very large (20+ GDT) differences. This is consistent with the results in Fig. 5a in which the depth of the MSA is relatively unimportant until it approaches a threshold value of around 30 sequences when the MSA size effects become quite large. We observe mostly overlapping effects between inclusion of BFD and MGnify, but having at least one of these metagenomics databases is very important for target classes that are poorly represented in UniRef, and having both was necessary to achieve full CASP accuracy.

Training regimen

To train, we use structures from the PDB with a maximum release date of 30 April 2018. Chains are sampled in inverse proportion to cluster size of a 40% sequence identity clustering. We then randomly crop them to 256 residues and assemble into batches of size 128. We train the model on Tensor Processing Unit (TPU) v3 with a batch size of 1 per TPU core, hence the model uses 128 TPU v3 cores. The model is trained until convergence (around 10 million samples) and further fine-tuned using longer crops of 384 residues, larger MSA stack and reduced learning rate (see Supplementary Methods 1.11 for the exact configuration). The initial training stage takes approximately 1 week, and the fine-tuning stage takes approximately 4 additional days.

The network is supervised by the FAPE loss and a number of auxiliary losses. First, the final pair representation is linearly projected to a binned distance distribution (distogram) prediction, scored with a cross-entropy loss. Second, we use random masking on the input MSAs and require the network to reconstruct the masked regions from the output MSA representation using a BERT-like loss³⁷. Third, the output single representations of the structure module are used to predict binned per-residue IDDT-Ca values. Finally, we use an auxiliary side-chain loss during training, and an auxiliary structure violation loss during fine-tuning. Detailed descriptions and weighting are provided in the Supplementary Information.

An initial model trained with the above objectives was used to make structure predictions for a Uniclust dataset of 355,993 sequences with the full MSAs. These predictions were then used to train a final model with identical hyperparameters, except for sampling examples 75% of the time from the Uniclust prediction set, with sub-sampled MSAs, and 25% of the time from the clustered PDB set.

We train five different models using different random seeds, some with templates and some without, to encourage diversity in the predictions (see Supplementary Table 5 and Supplementary Methods 1.12.1

for details). We also fine-tuned these models after CASP14 to add a pTM prediction objective (Supplementary Methods 1.9.7) and use the obtained models for Fig. 2d.

Inference regimen

We inference the five trained models and use the predicted confidence score to select the best model per target.

Using our CASP14 configuration for AlphaFold, the trunk of the network is run multiple times with different random choices for the MSA cluster centres (see Supplementary Methods 1.11.2 for details of the ensembling procedure). The full time to make a structure prediction varies considerably depending on the length of the protein. Representative timings for the neural network using a single model on V100 GPU are 4.8 min with 256 residues, 9.2 min with 384 residues and 18 h at 2,500 residues. These timings are measured using our open-source code, and the open-source code is notably faster than the version we ran in CASP14 as we now use the XLA compiler⁷⁵.

Since CASP14, we have found that the accuracy of the network without ensembling is very close or equal to the accuracy with ensembling and we turn off ensembling for most inference. Without ensembling, the network is 8× faster and the representative timings for a single model are 0.6 min with 256 residues, 1.1 min with 384 residues and 2.1 h with 2,500 residues.

Inferencing large proteins can easily exceed the memory of a single GPU. For a V100 with 16 GB of memory, we can predict the structure of proteins up to around 1,300 residues without ensembling and the 256- and 384-residue inference times are using the memory of a single GPU. The memory usage is approximately quadratic in the number of residues, so a 2,500-residue protein involves using unified memory so that we can greatly exceed the memory of a single V100. In our cloud setup, a single V100 is used for computation on a 2,500-residue protein but we requested four GPUs to have sufficient memory.

Searching genetic sequence databases to prepare inputs and final relaxation of the structures take additional central processing unit (CPU) time but do not require a GPU or TPU.

Metrics

The predicted structure is compared to the true structure from the PDB in terms of IDDT metric³⁴, as this metric reports the domain accuracy without requiring a domain segmentation of chain structures. The distances are either computed between all heavy atoms (IDDT) or only the C α atoms to measure the backbone accuracy (IDDT-C α). As IDDT-C α only focuses on the C α atoms, it does not include the penalty for structural violations and clashes. Domain accuracies in CASP are reported as GDT³³ and the TM-score²⁷ is used as a full chain global superposition metric.

We also report accuracies using the r.m.s.d.₉₅ (C α r.m.s.d. at 95% coverage). We perform five iterations of (1) a least-squares alignment of the predicted structure and the PDB structure on the currently chosen C α atoms (using all C α atoms in the first iteration); (2) selecting the 95% of C α atoms with the lowest alignment error. The r.m.s.d. of the atoms chosen for the final iterations is the r.m.s.d.₉₅. This metric is more robust to apparent errors that can originate from crystal structure artefacts, although in some cases the removed 5% of residues will contain genuine modelling errors.

Test set of recent PDB sequences

For evaluation on recent PDB sequences (Figs. 2a–d, 4a, 5a), we used a copy of the PDB downloaded 15 February 2021. Structures were filtered to those with a release date after 30 April 2018 (the date limit for inclusion in the training set for AlphaFold). Chains were further filtered to remove sequences that consisted of a single amino acid as well as sequences with an ambiguous chemical component at any residue position. Exact duplicates were removed, with the chain with the most resolved C α atoms used as the representative sequence. Subsequently,

structures with less than 16 resolved residues, with unknown residues or solved by NMR methods were removed. As the PDB contains many near-duplicate sequences, the chain with the highest resolution was selected from each cluster in the PDB 40% sequence clustering of the data. Furthermore, we removed all sequences for which fewer than 80 amino acids had the alpha carbon resolved and removed chains with more than 1,400 residues. The final dataset contained 10,795 protein sequences.

The procedure for filtering the recent PDB dataset based on prior template identity was as follows. Hmsearch was run with default parameters against a copy of the PDB SEQRES fasta downloaded 15 February 2021. Template hits were accepted if the associated structure had a release date earlier than 30 April 2018. Each residue position in a query sequence was assigned the maximum identity of any template hit covering that position. Filtering then proceeded as described in the individual figure legends, based on a combination of maximum identity and sequence coverage.

The MSA depth analysis was based on computing the normalized number of effective sequences (N_{eff}) for each position of a query sequence. Per-residue N_{eff} values were obtained by counting the number of non-gap residues in the MSA for this position and weighting the sequences using the N_{eff} scheme⁷⁶ with a threshold of 80% sequence identity measured on the region that is non-gap in either sequence.

Reporting summary

Further information on research design is available in the Nature Research Reporting Summary linked to this paper.

Data availability

All input data are freely available from public sources.

Structures from the PDB were used for training and as templates (<https://www.wwpdb.org/ftp/pdb-ftp-sites>; for the associated sequence data and 40% sequence clustering see also https://ftp.wwpdb.org/pub/pdb/derived_data/ and <https://cdn.rcsb.org/resources/sequence/clusters/bc-40.out>). Training used a version of the PDB downloaded 28 August 2019, while the CASP14 template search used a version downloaded 14 May 2020. The template search also used the PDB70 database, downloaded 13 May 2020 (https://wwwuser.gwdg.de/~compbiol/data/hhsuite/databases/hhsuite_dbs/).

We show experimental structures from the PDB with accession numbers 6Y4F⁷⁷, 6YJ1⁷⁸, 6VR4⁷⁹, 6SK0⁸⁰, 6FES⁸¹, 6W6W⁸², 6T1Z⁸³ and 7JTL⁸⁴.

For MSA lookup at both the training and prediction time, we used UniRef90 v.2020_01 (https://ftp.ebi.ac.uk/pub/databases/uniprot/previous_releases/release-2020_01/uniref/), BFD (<https://bfd.mmseqs.com>), Uniclust30 v.2018_08 (https://wwwuser.gwdg.de/~compbiol/uniclust/2018_08/) and MGnify clusters v.2018_12 (https://ftp.ebi.ac.uk/pub/databases/metagenomics/peptide_database/2018_12/). Uniclust30 v.2018_08 was also used as input for constructing a distillation structure dataset.

Code availability

Source code for the AlphaFold model, trained weights and inference script are available under an open-source license at <https://github.com/deepmind/alphafold>.

Neural networks were developed with TensorFlow v.1 (<https://github.com/tensorflow/tensorflow>), Sonnet v.1 (<https://github.com/deepmind/sonnet>), JAX v.0.1.69 (<https://github.com/google/jax/>) and Haiku v.0.0.4 (<https://github.com/deepmind/dm-haiku>). The XLA compiler is bundled with JAX and does not have a separate version number.

For MSA search on BFD+Uniclust30, and for template search against PDB70, we used HHBlits and HHSearch from hh-suite v.3.0-beta.3 release 14/07/2017 (<https://github.com/soedinglab/hh-suite>). For MSA search on UniRef90 and clustered MGnify, we used jackhmmer from

HMMER v.3.3 (<http://eddylib.org/software/hmmer/>). For constrained relaxation of structures, we used OpenMM v.7.3.1 (<https://github.com/openmm/openmm>) with the Amber99sb force field.

Construction of BFD used MMseqs2 v.925AF (<https://github.com/soedinglab/MMseqs2>) and FAMSA v.1.2.5 (<https://github.com/refresh-bio/FAMSA>).

Data analysis used Python v.3.6 (<https://www.python.org/>), NumPy v.1.16.4 (<https://github.com/numpy/numpy>), SciPy v.1.2.1 (<https://www.scipy.org/>), seaborn v.0.11.1 (<https://github.com/mwaskom/seaborn>), Matplotlib v.3.3.4 (<https://github.com/matplotlib/matplotlib>), bokeh v.1.4.0 (<https://github.com/bokeh/bokeh>), pandas v.1.1.5 (<https://github.com/pandas-dev/pandas>), plotnine v.0.8.0 (<https://github.com/has2k1/plotnine>), statsmodels v.0.12.2 (<https://github.com/statsmodels/statsmodels>) and Colab (<https://research.google.com/colaboratory>). TM-align v.20190822 (<https://zhanglab.dcm.med.umich.edu/TM-align/>) was used for computing TM-scores. Structure visualizations were created in PyMol v.2.3.0 (<https://github.com/schrodinger/pymol-open-source>).

59. Ingraham, J., Garg, V. K., Barzilay, R. & Jaakkola, T. Generative models for graph-based protein design. in *Proc. 33rd Conference on Neural Information Processing Systems* (2019).
60. Johnson, L. S., Eddy, S. R. & Portugaly, E. Hidden Markov model speed heuristic and iterative HMM search procedure. *BMC Bioinformatics* **11**, 431 (2010).
61. Remmert, M., Biegert, A., Hauser, A. & Söding, J. HHblits: lightning-fast iterative protein sequence searching by HMM-HMM alignment. *Nat. Methods* **9**, 173–175 (2012).
62. The UniProt Consortium. UniProt: the universal protein knowledgebase in 2021. *Nucleic Acids Res.* **49**, D480–D489 (2020).
63. Steinegger, M. & Söding, J. Clustering huge protein sequence sets in linear time. *Nat. Commun.* **9**, 2542 (2018).
64. Steinegger, M. & Söding, J. MMseqs2 enables sensitive protein sequence searching for the analysis of massive data sets. *Nat. Biotechnol.* **35**, 1026–1028 (2017).
65. Deorowicz, S., Debudaj-Grabysz, A. & Gudyś, A. FAMSA: fast and accurate multiple sequence alignment of huge protein families. *Sci. Rep.* **6**, 33964 (2016).
66. Steinegger, M. et al. HH-suite3 for fast remote homology detection and deep protein annotation. *BMC Bioinformatics* **20**, 473 (2019).
67. Suzek, B. E., Wang, Y., Huang, H., McGarvey, P. B. & Wu, C. H. UniRef clusters: a comprehensive and scalable alternative for improving sequence similarity searches. *Bioinformatics* **31**, 926–932 (2015).
68. Eddy, S. R. Accelerated profile HMM searches. *PLOS Comput. Biol.* **7**, e1002195 (2011).
69. Eastman, P. et al. OpenMM 7: rapid development of high performance algorithms for molecular dynamics. *PLOS Comput. Biol.* **13**, e1005659 (2017).
70. Ashish, A. M. A. et al. TensorFlow: large-scale machine learning on heterogeneous systems. Preprint at <https://arxiv.org/abs/1603.04467> (2015).
71. Reynolds, M. et al. Open sourcing Sonnet – a new library for constructing neural networks. DeepMind <https://deepmind.com/blog/open-sourcing-sonnet/> (7 April 2017).
72. Harris, C. R. et al. Array programming with NumPy. *Nature* **585**, 357–362 (2020).
73. Van Rossum, G. & Drake, F. L. Python 3 Reference Manual (CreateSpace, 2009).
74. Bisong, E. in *Building Machine Learning and Deep Learning Models on Google Cloud Platform: A Comprehensive Guide for Beginners* 59–64 (Apress, 2019).
75. TensorFlow. XLA: Optimizing Compiler for TensorFlow. <https://www.tensorflow.org/xla> (2018).

76. Wu, T., Hou, J., Adhikari, B. & Cheng, J. Analysis of several key factors influencing deep learning-based inter-residue contact prediction. *Bioinformatics* **36**, 1091–1098 (2020).
77. Jiang, W. et al. MrpH, a new class of metal-binding adhesin, requires zinc to mediate biofilm formation. *PLoS Pathog.* **16**, e1008707 (2020).
78. Dunne, M., Ernst, P., Sobieraj, A., Pluckthun, A. & Loessner, M. J. The M23 peptidase domain of the Staphylococcal phage 2638A endolysin. PDB <https://doi.org/10.2210/pdb6YJ1/pdb> (2020).
79. Drobysheva, A. V. et al. Structure and function of virion RNA polymerase of a crAss-like phage. *Nature* **589**, 306–309 (2021).
80. Flaughnatti, N. et al. Structural basis for loading and inhibition of a bacterial T6SS phospholipase effector by the VgrG spike. *EMBO J.* **39**, e104129 (2020).
81. ElGamacy, M. et al. An interface-driven design strategy yields a novel, corrugated protein architecture. *ACS Synth. Biol.* **7**, 2226–2235 (2018).
82. Lim, C. J. et al. The structure of human CST reveals a decameric assembly bound to telomeric DNA. *Science* **368**, 1081–1085 (2020).
83. Debruycker, V. et al. An embedded lipid in the multidrug transporter LmrP suggests a mechanism for polyspecificity. *Nat. Struct. Mol. Biol.* **27**, 829–835 (2020).
84. Flower, T. G. et al. Structure of SARS-CoV-2 ORF8, a rapidly evolving immune evasion protein. *Proc. Natl Acad. Sci. USA* **118**, e2021785118 (2021).

Acknowledgements We thank A. Rustemi, A. Gu, A. Guseynov, B. Hechtman, C. Beattie, C. Jones, C. Donner, E. Parisotto, E. Elsen, F. Popovici, G. Necula, H. Maclean, J. Menick, J. Kirkpatrick, J. Molloy, J. Yim, J. Stanway, K. Simonyan, L. Sifre, L. Martens, M. Johnson, M. O'Neill, N. Antropova, R. Hadsell, S. Blackwell, S. Das, S. Hou, S. Gouws, S. Wheelwright, T. Hennigan, T. Ward, Z. Wu, Z. Avsec and the Research Platform Team for their contributions; M. Mirdita for his help with the datasets; M. Piovesan-Forster, A. Nelson and R. Kemp for their help managing the project; the JAX, TensorFlow and XLA teams for detailed support and enabling machine learning models of the complexity of AlphaFold; our colleagues at DeepMind, Google and Alphabet for their encouragement and support; and J. Moulton and the CASP14 organizers, and the experimentalists whose structures enabled the assessment. M.S. acknowledges support from the National Research Foundation of Korea grant (2019R1A6A1A10073437, 2020M3A9G7103933) and the Creative-Pioneering Researchers Program through Seoul National University.

Author contributions J.J. and D.H. led the research. J.J., R.E., A. Pritzel, M.F., O.R., R.B., A. Potapenko, S.A.A.K., B.R.-P., J.A., M.P., T. Berghammer and O.V. developed the neural network architecture and training. T.G., A.Z., K.T., R.B., A.B., R.E., A.J.B., A.C., S.N., R.J., D.R., M.Z. and S.B. developed the data, analytics and inference systems. D.H., K.K., P.K., C.M. and E.C. managed the research. T.G. led the technical platform. P.K., A.W.S., K.K., O.V., D.S., S.P. and T. Back contributed technical advice and ideas. M.S. created the BFD genomics database and provided technical assistance on HHblits. D.H., R.E., A.W.S. and K.K. conceived the AlphaFold project. J.J., R.E. and A.W.S. conceived the end-to-end approach. J.J., A. Pritzel, O.R., A. Potapenko, R.E., M.F., T.G., K.T., C.M. and D.H. wrote the paper.

Competing interests J.J., R.E., A. Pritzel, T.G., M.F., O.R., R.B., A.B., S.A.A.K., D.R. and A.W.S. have filed non-provisional patent applications 16/701,070 and PCT/EP2020/084238, and provisional patent applications 63/107,362, 63/118,917, 63/118,918, 63/118,921 and 63/118,919, each in the name of DeepMind Technologies Limited, each pending, relating to machine learning for predicting protein structures. The other authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41586-021-03819-2>.

Correspondence and requests for materials should be addressed to J.J. or D.H.

Peer review information Nature thanks Mohammed AlQuraishi, Charlotte Deane and Yang Zhang for their contribution to the peer review of this work.

Reprints and permissions information is available at <http://www.nature.com/reprints>.

Reporting Summary

Nature Research wishes to improve the reproducibility of the work that we publish. This form provides structure for consistency and transparency in reporting. For further information on Nature Research policies, see our [Editorial Policies](#) and the [Editorial Policy Checklist](#).

Statistics

For all statistical analyses, confirm that the following items are present in the figure legend, table legend, main text, or Methods section.

n/a Confirmed

- ☒ ☒ The exact sample size (n) for each experimental group/condition, given as a discrete number and unit of measurement
- ☒ ☐ A statement on whether measurements were taken from distinct samples or whether the same sample was measured repeatedly
- ☐ ☒ The statistical test(s) used AND whether they are one- or two-sided
Only common tests should be described solely by name; describe more complex techniques in the Methods section.
- ☒ ☐ A description of all covariates tested
- ☐ ☒ A description of any assumptions or corrections, such as tests of normality and adjustment for multiple comparisons
- ☐ ☒ A full description of the statistical parameters including central tendency (e.g. means) or other basic estimates (e.g. regression coefficient) AND variation (e.g. standard deviation) or associated estimates of uncertainty (e.g. confidence intervals)
- ☒ ☐ For null hypothesis testing, the test statistic (e.g. F, t, r) with confidence intervals, effect sizes, degrees of freedom and P value noted
Give P values as exact values whenever suitable.
- ☐ ☒ For Bayesian analysis, information on the choice of priors and Markov chain Monte Carlo settings
- ☒ ☐ For hierarchical and complex designs, identification of the appropriate level for tests and full reporting of outcomes
- ☐ ☒ Estimates of effect sizes (e.g. Cohen's d, Pearson's r), indicating how they were calculated

Our web collection on [statistics for biologists](#) contains articles on many of the points above.

Software and code

Policy information about [availability of computer code](#)

Data collection

Source code for the AlphaFold model, trained weights, and inference script will be made available under an open-source license at <https://github.com/deepmind/> upon publication.

Neural networks were developed with TensorFlow v1 (<https://github.com/tensorflow/tensorflow>), Sonnet v1 (<https://github.com/deepmind/sonnet>), JAX v0.1.69 (<https://github.com/google/jax>), and Haiku v0.0.4 (<https://github.com/deepmind/dm-haiku>). The XLA compiler is bundled with JAX and does not have a separate version number.

For MSA search on BFD+Uniclust30, and for template search against PDB70, we used HHBlits and HHSearch from hh-suite v3.0-beta.3 14/07/2017 (<https://github.com/soedinglab/hh-suite>). For MSA search on UniRef90 and clustered MGnify, we used jackhmmer from HMMER v3.3 (<http://eddylib.org/software/hmmer/>). For constrained relaxation of structures, we used OpenMM v7.3.1 (<https://github.com/openmm/openmm>) with the Amber99sb force field.

Construction of BFD used MMseqs2 version 925AF (<https://github.com/soedinglab/MMseqs2>) and FAMSA v1.2.5 (<https://github.com/refresh-bio/FAMSA>).

Data analysis

Data analysis used Python v3.6 (<https://www.python.org/>), NumPy v1.16.4 (<https://github.com/numpy/numpy>), SciPy v1.2.1 (<https://www.scipy.org/>), seaborn v0.11.1 (<https://github.com/mwaskom/seaborn>), Matplotlib v3.3.4 (<https://github.com/matplotlib/matplotlib>), bokeh v1.4.0 (<https://github.com/bokeh/bokeh>), pandas v1.1.5 (<https://github.com/pandas-dev/pandas>), plotnine v0.8.0 (<https://github.com/has2k1/plotnine>), statsmodels v0.12.2 (<https://github.com/statsmodels/statsmodels>) and Colab (<https://research.google.com/colaboratory>). TM-align v20190822 (<https://zhanglab.dcm.med.umich.edu/TM-align/>) was used for computing TM-scores. Structure visualizations were created in Pymol v2.3.0 (<https://github.com/schrodinger/pymol-open-source>).

For manuscripts utilizing custom algorithms or software that are central to the research but not yet described in published literature, software must be made available to editors and reviewers. We strongly encourage code deposition in a community repository (e.g. GitHub). See the Nature Research [guidelines for submitting code & software](#) for further information.

Data

Policy information about availability of data

All manuscripts must include a **data availability statement**. This statement should provide the following information, where applicable:

- Accession codes, unique identifiers, or web links for publicly available datasets
- A list of figures that have associated raw data
- A description of any restrictions on data availability

All input data are freely available from public sources.

Structures from the PDB were used for training and as templates (<https://www.wwpdb.org/ftp/pdb-ftp-sites>; for the associated sequence data and 40% sequence clustering see also https://ftp.wwpdb.org/pub/pdb/derived_data/ and <https://cdn.rcsb.org/resources/sequence/clusters/bc-40.out>). Training used a version of the PDB downloaded 28/08/2019, while CASP14 template search used a version downloaded 14/05/2020. Template search also used the PDB70 database, downloaded 13/05/2020 (https://wwwuser.gwdg.de/~compbiol/data/hhsuite/databases/hhsuite_dbs/).

We show experimental structures from the PDB with accessions 6Y4F77, 6YJ178, 6VR479, 6SK080, 6FES81, 6W6W82, 6T1Z83, and 7JTL84.

For MSA lookup at both training and prediction time, we used UniRef90 v2020_01 (https://ftp.ebi.ac.uk/pub/databases/uniprot/previous_releases/release-2020_01/uniref/), BFD (<https://bfd.mmseqs.com>), Uniclust30 v2018_08 (https://wwwuser.gwdg.de/~compbiol/uniclust/2018_08/), and MGnify clusters v2018_12 (https://ftp.ebi.ac.uk/pub/databases/metagenomics/peptide_database/2018_12/). Uniclust30 v2018_08 was further used as input for constructing a distillation structure dataset.

Field-specific reporting

Please select the one below that is the best fit for your research. If you are not sure, read the appropriate sections before making your selection.

☒ Life sciences ☐ Behavioural & social sciences ☐ Ecological, evolutionary & environmental sciences

For a reference copy of the document with all sections, see [nature.com/documents/nr-reporting-summary-flat.pdf](https://www.nature.com/documents/nr-reporting-summary-flat.pdf)

Life sciences study design

All studies must disclose on these points even when the disclosure is negative.

Sample size	No sample size was chosen; the method was evaluated on the full CASP14 benchmark set, and all PDB chains not in the training set (subject to the exclusions noted below).
Data exclusions	The recent PDB set was filtered (see Methods for full details). Briefly this excludes chains with too few resolved residues, longer than 1400 residues, solved by NMR or with unknown/ambiguous residues. This set was also redundancy reduced (by taking representatives from a sequence clustering), and for some figures a sequence similarity-based filter was applied to remove entries too similar to the training set (see Methods and figure legends for details).
Replication	Not applicable, no experimental work is described in this study. The results are the output of a computational method which will be made available.
Randomization	Not applicable, we are not making a comparison between two groups
Blinding	Not applicable, we are not making a comparison between two groups

Reporting for specific materials, systems and methods

We require information from authors about some types of materials, experimental systems and methods used in many studies. Here, indicate whether each material, system or method listed is relevant to your study. If you are not sure if a list item applies to your research, read the appropriate section before selecting a response.


Materials & experimental systems

n/a	Involved in the study
☒	☐ Antibodies
☒	☐ Eukaryotic cell lines
☒	☐ Palaeontology and archaeology
☒	☐ Animals and other organisms
☒	☐ Human research participants
☒	☐ Clinical data
☒	☐ Dual use research of concern

Methods

n/a	Involved in the study
☒	☐ ChIP-seq
☒	☐ Flow cytometry
☒	☐ MRI-based neuroimaging

AI Content



	Text coverage	Words
☐ AI Text	0%	0
☐ Human Text	100%	12,076
Excluded		
☐ Omitted Words		0

About AI Detection

Our AI Detector is the only enterprise-level solution that can verify if the content was written by a human or generated by AI, including source code and text that has been plagiarized or modified. [Learn more](#)

☐ AI Text

A body of text that has been generated or altered by AI technology. [Learn more](#)

☐ Human Text

Any text that has been fully written by a human and has not been altered or generated by AI. [Learn more](#)

Copyleaks AI Detector Effectiveness

Credible data at scale, coupled with machine learning and widespread adoption, allows us to continually refine and improve our ability to understand complex text patterns, resulting in over 99% accuracy—far higher than any other AI detector—and improving daily. [Learn more](#)

Ideal Text Length

The higher the character count, the easier for our technology to determine irregular patterns, which results in a higher confidence rating for AI detection. [Learn more](#)

Reasons It Might Be AI When You Think It's Not

The AI Detector can detect a variety of AI-generated text, including tools that use AI technology to paraphrase content, auto-complete sentences, and more. [Learn more](#)

User AI Alert History

Historical data of how many times a user has been flagged for potentially having AI text within their content. [Learn more](#)

AI Logic

The number of times a phrase was found more frequently in AI vs human text is shown according to low, medium, and high frequency. [Learn more](#)