

Machine Learning for Natural Science

Joshua Yao-Yu Lin (林曜宇)
[2025-12-12 @ACML Tutorial]

Joshua Yao-Yu Lin



- Currently a Machine Learning Scientist at Genentech (New York office)
- UIUC Physics Ph.D. (2016-2022), MS at NTU, and BS at NTHU in Taiwan.
- My past research spans a wide range of Machine Learning application for astrophysics, including black hole image and dark matter/strong lensing
- ML Research Interest: **ML for Natural Science, ML for drug discovery & Protein Design**
- Previous ML experience: Genentech Postdoc with Prof. Kyunghyun Cho, Simons Foundation/Flatiron Institute (CCA), Google Research (2021)
- I like: Traveling, Jazz, Bouldering/Climbing, Brewing hard cider

Prescient Design@Genentech/Roche



**Vladimir
Gligorijevic**

Co-Founder and Senior
Director, Prescient
Design, Genentech



Richard Bonneau

Co-Founder and
Executive Director,
Prescient Design,
Genentech



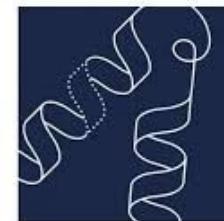
Kyunghyun Cho

Co-Founder and Senior
Director, Prescient
Design, Genentech

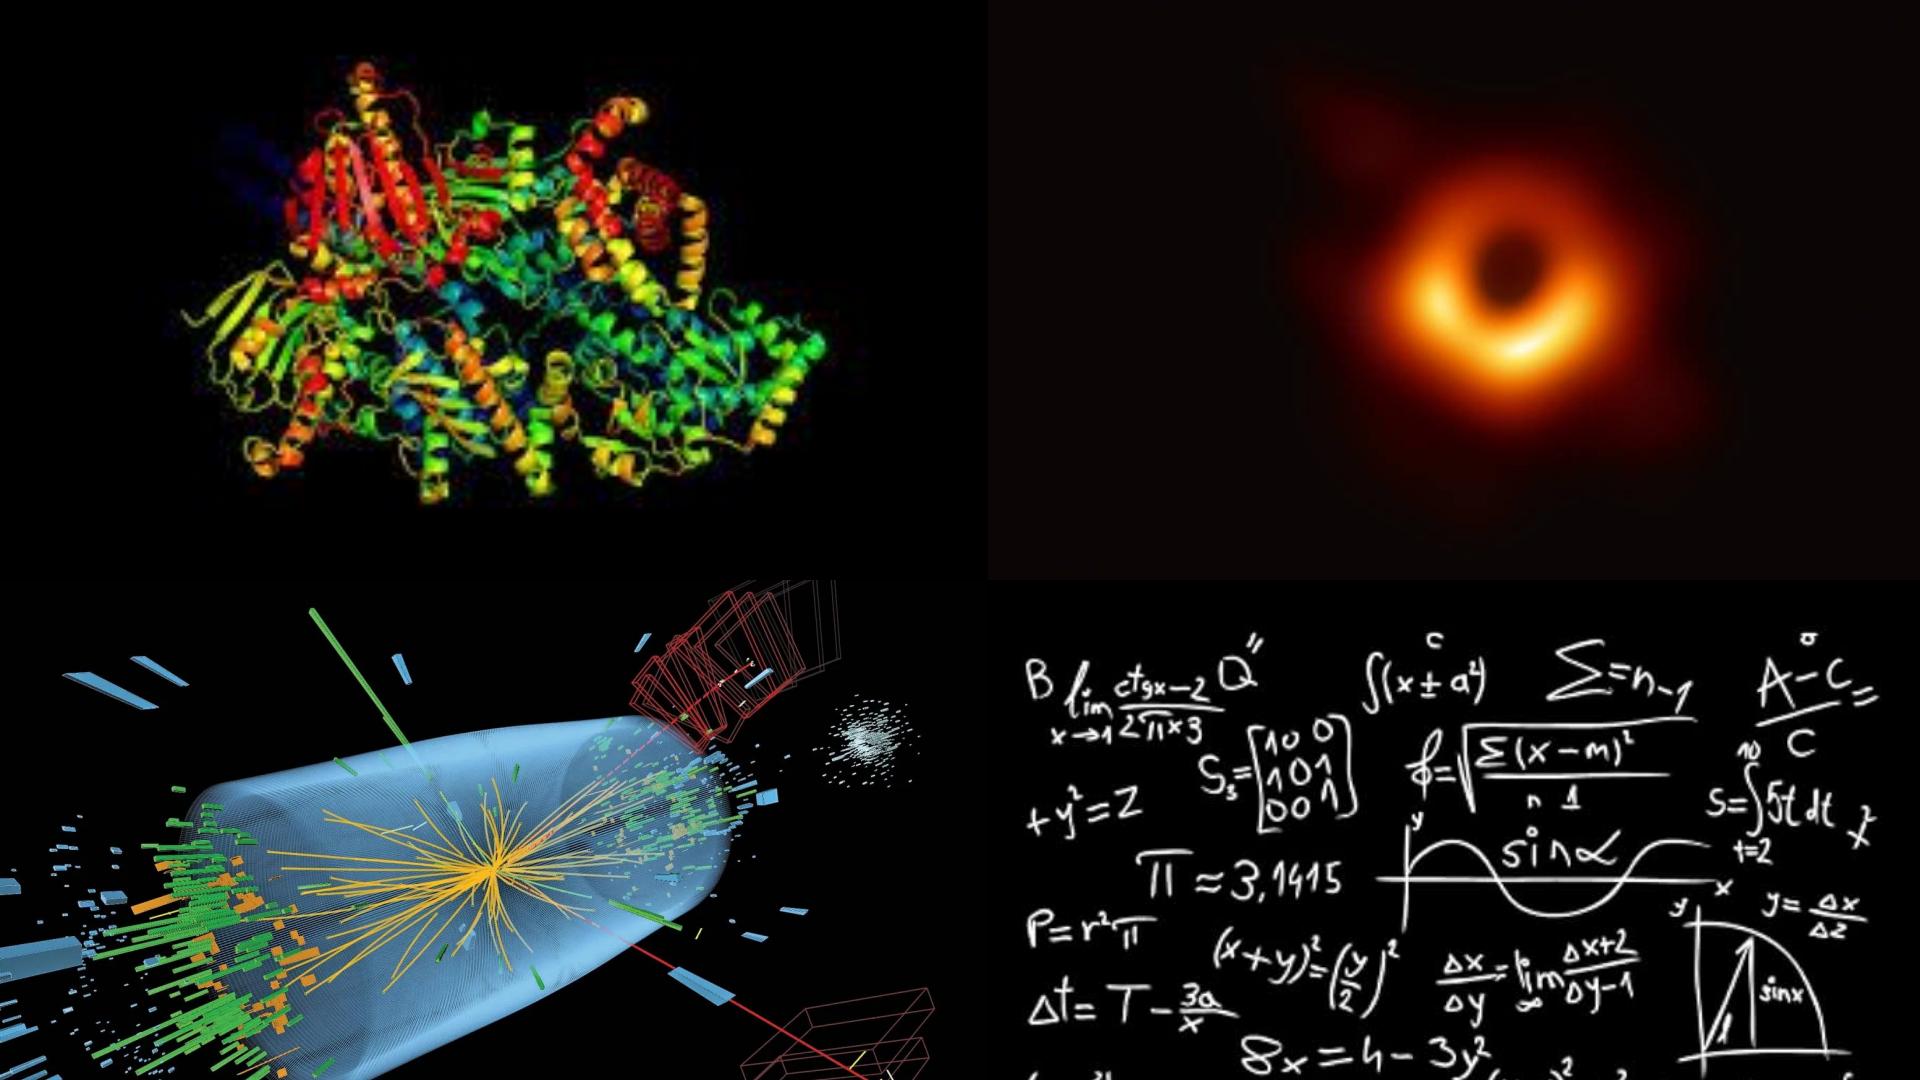
- Founded in Jan 2021, focusing on machine learning for **Protein Design/drug discovery**
- Acquired by **Genentech/Roche** ~ August 2021
- Around 70 people in the team (ML Scientist/Engineer, Bio/Chem)
- We're hiring!

Prev. @Flatiron Institute/NYU

NYU CS/Data Science



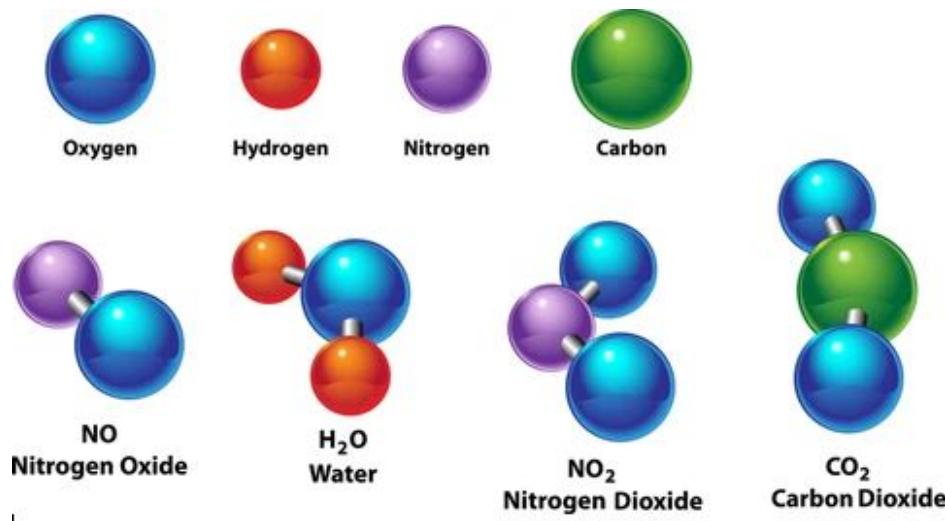
Prescient
Design
A Genentech Accelerator





GROUP		PERIOD																		18																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																									
1		2																		He																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																									
PERIOD	Alkali Metals	Non-metals	3		4		5		6		7		8		9		10		11		12		13		14		15		16		17		18																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																												
	Alkaline Earth Metals	Halogens	3		4		5		6		7		8		9		10		11		12		13		14		15		16		17		18																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																												
1	H	Hydrogen 3.99	Li	Be	Boron 8.004	12	Na	Mg	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35	36	37	38	39	40	41	42	43	44	45	46	47	48	49	50	51	52	53	54	55	56	57	58	59	60	61	62	63	64	65	66	67	68	69	70	71	72	73	74	75	76	77	78	79	80	81	82	83	84	85	86	87	88	89	90	91	92	93	94	95	96	97	98	99	100	101	102	103	104	105	106	107	108	109	110	111	112	113	114	115	116	117	118	119	120	121	122	123	124	125	126	127	128	129	130	131	132	133	134	135	136	137	138	139	140	141	142	143	144	145	146	147	148	149	150	151	152	153	154	155	156	157	158	159	160	161	162	163	164	165	166	167	168	169	170	171	172	173	174	175	176	177	178	179	180	181	182	183	184	185	186	187	188	189	190	191	192	193	194	195	196	197	198	199	200	201	202	203	204	205	206	207	208	209	210	211	212	213	214	215	216	217	218	219	220	221	222	223	224	225	226	227	228	229	230	231	232	233	234	235	236	237	238	239	240	241	242	243	244	245	246	247	248	249	250	251	252	253	254	255	256	257	258	259	260	261	262	263	264	265	266	267	268	269	270	271	272	273	274	275	276	277	278	279	280	281	282	283	284	285	286	287	288	289	290	291	292	293	294	295	296	297	298	299	300	301	302	303	304	305	306	307	308	309	310	311	312	313	314	315	316	317	318	319	320	321	322	323	324	325	326	327	328	329	330	331	332	333	334	335	336	337	338	339	340	341	342	343	344	345	346	347	348	349	350	351	352	353	354	355	356	357	358	359	360	361	362	363	364	365	366	367	368	369	370	371	372	373	374	375	376	377	378	379	380	381	382	383	384	385	386	387	388	389	390	391	392	393	394	395	396	397	398	399	400	401	402	403	404	405	406	407	408	409	410	411	412	413	414	415	416	417	418	419	420	421	422	423	424	425	426	427	428	429	430	431	432	433	434	435	436	437	438	439	440	441	442	443	444	445	446	447	448	449	450	451	452	453	454	455	456	457	458	459	460	461	462	463	464	465	466	467	468	469	470	471	472	473	474	475	476	477	478	479	480	481	482	483	484	485	486	487	488	489	490	491	492	493	494	495	496	497	498	499	500	501	502	503	504	505	506	507	508	509	510	511	512	513	514	515	516	517	518	519	520	521	522	523	524	525	526	527	528	529	530	531	532	533	534	535	536	537	538	539	540	541	542	543	544	545	546	547	548	549	550	551	552	553	554	555	556	557	558	559	560	561	562	563	564	565	566	567	568	569	570	571	572	573	574	575	576	577	578	579	580	581	582	583	584	585	586	587	588	589	590	591	592	593	594	595	596	597	598	599	600	601	602	603	604	605	606	607	608	609	610	611	612	613	614	615	616	617	618	619	620	621	622	623	624	625	626	627	628	629	630	631	632	633	634	635	636	637	638	639	640	641	642	643	644	645	646	647	648	649	650	651	652	653	654	655	656	657	658	659	660	661	662	663	664	665	666	667	668	669	670	671	672	673	674	675	676	677	678	679	680	681	682	683	684	685	686	687	688	689	690	691	692	693	694	695	696	697	698	699	700	701	702	703	704	705	706	707	708	709	710	711	712	713	714	715	716	717	718	719	720	721	722	723	724	725	726	727	728	729	730	731	732	733	734	735	736	737	738	739	740	741	742	743	744	745	746	747	748	749	750	751	752	753	754	755	756	757	758	759	760	761	762	763	764	765	766	767	768	769	770	771	772	773	774	775	776	777	778	779	780	781	782	783	784	785	786	787	788	789	790	791	792	793	794	795	796	797	798	799	800	801	802	803	804	805	806	807	808	809	8010	8011	8012	8013	8014	8015	8016	8017	8018	8019	8020	8021	8022	8023	8024	8025	8026	8027	8028	8029	8030	8031	8032	8033	8034	8035	8036	8037	8038	8039	8040	8041	8042	8043	8044	8045	8046	8047	8048	8049	8050	8051	8052	8053	8054	8055	8056	8057	8058	8059	8060	8061	8062	8063	8064	8065	8066	8067	8068	8069	8070	8071	8072	8073	8074	8075	8076	8077	8078	8079	8080	8081	8082	8083	8084	8085	8086	8087	8088	8089	8090	8091	8092	8093	8094	8095	8096	8097	8098	8099	80100	80101	80102	80103	80104	80105	80106	80107	80108	80109	80110	80111	80112	80113	80114	80115	80116	80117	80118	80119	80120	80121	80122	80123	80124	80125	80126	80127	80128	80129	80130	80131	80132	80133	80134	80135	80136	80137	80138	80139	80140	80141	80142	80143	80144	80145	80146	80147	80148	80149	80150	80151	80152	80153	80154	80155	80156	80157	80158	80159	80160	80161	80162	80163	80164	80165	80166	80167	80168	80169	80170	80171	80172	80173	80174	80175	80176	80177	80178	80179	80180	80181	80182	80183	80184	80185	80186	80187	80188	80189	80190	80191	80192	80193	80194	80195	80196	80197	80198	80199	80200	80201	80202	80203	80204	80205	80206	80207	80208	80209	80210	80211	80212	80213	80214	80215	80216	80217	80218	80219	80220	80221	80222	80223	80224	80225	80226	80227	80228	80229	80230	80231	80232	80233	80234	80235	80236	80237	80238	80239	80240	80241	80242	80243	80244	80245	80246	80247	80248	80249	80250	80251	80252	80253	80254	80255	80256	80257	80258	80259	80260	80261	80262	80263	80264	80265	80266	80267	80268	80269	80270	80271	80272	80273	80274	80275	80276	80277	80278	80279	80280	80281	80282	80283	80284	80285	80286	80287	80288	80289	80290	80291	80292	80293	80294	80295	80296	80297	80298	80299	80300	80301	80302	80303	80304	80305	80306	80307	80308	80309	80310	80311	80312	80313	80314	80315	80316	80317	80318	80319	80320	80321	80322	80323	80324	80325	80326	80327	80328	80329	80330	80331	80332	80333	80334	80335	80336	80337	80338	80339	80340	80341	80342	80343	80344	80345	80346	80347	80348	80349	80350	80351	80352	80353	80354	80355	80356	80357	80358	80359	80360	80361	80362	80363	80364	80365	80366	80367	80368	80369	80370	80371	80372	80373	80374	80375	80376	80377	80378	80379	80380	80381	80382	80383	80384	80385	80386	80387	80388	80389	80390	80391	80392	80393	80394	80395	80396	80397	80398	80399	80400	80401	80402	80403	80404	80405	80406	80407	80408	80409	80410	80411	80412	80413	80414	80415	80416	80417	80418	80419	80420	80421	80422	80423	80424	80425	80426	80427	80428	80429	80430	80431	80432	80433

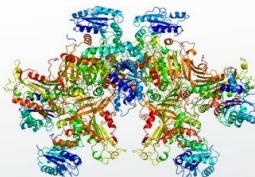
Atom -> Molecules



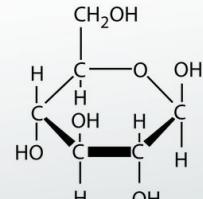
shutterstock.com • 1142362373

Large Molecules

LESSON SUMMARY



proteins



carbohydrates

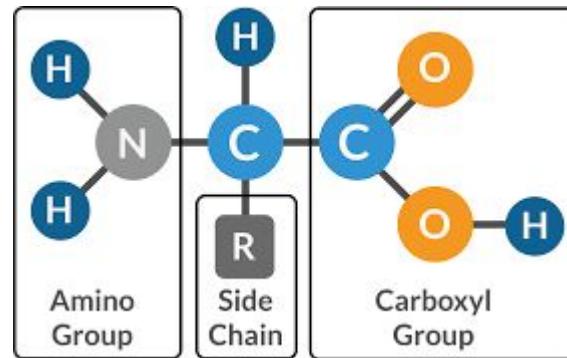


lipids

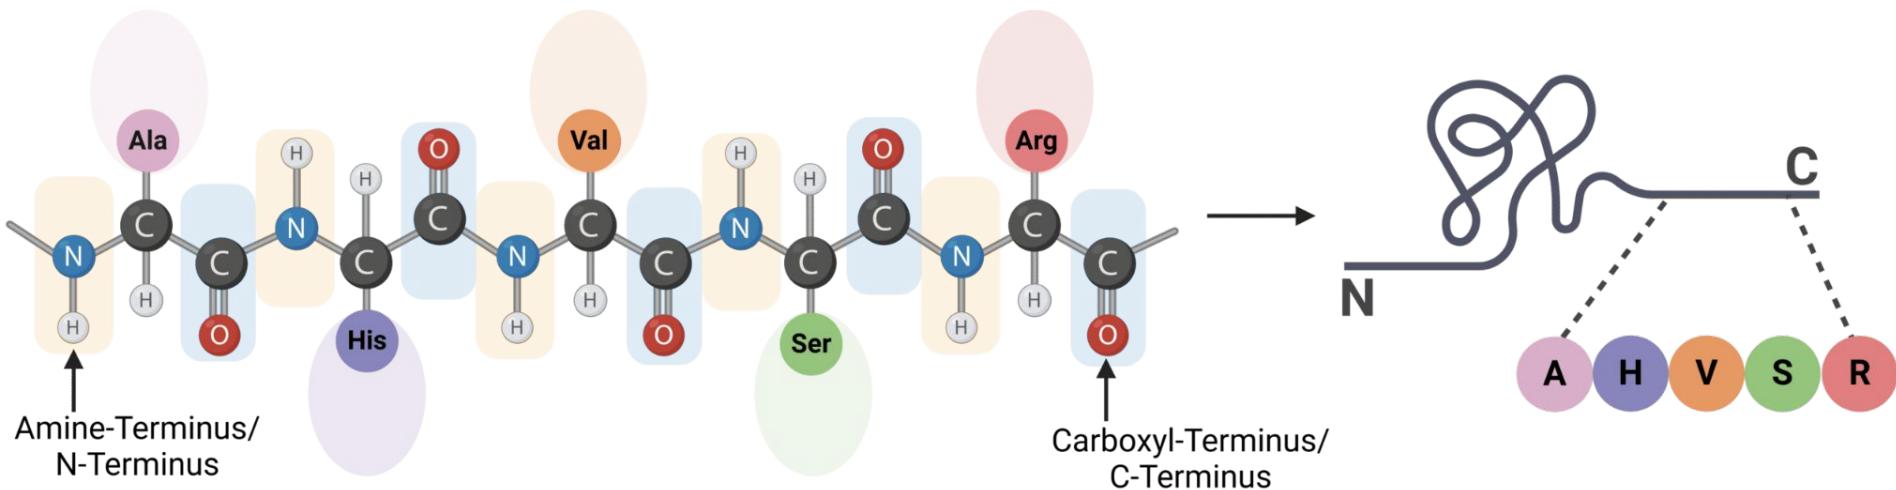


nucleic acids

Amino Acids



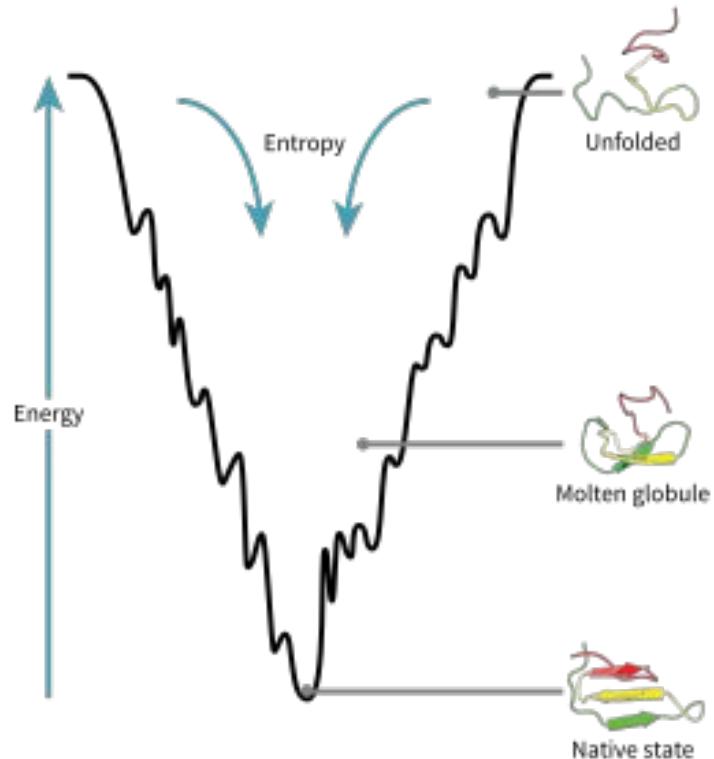
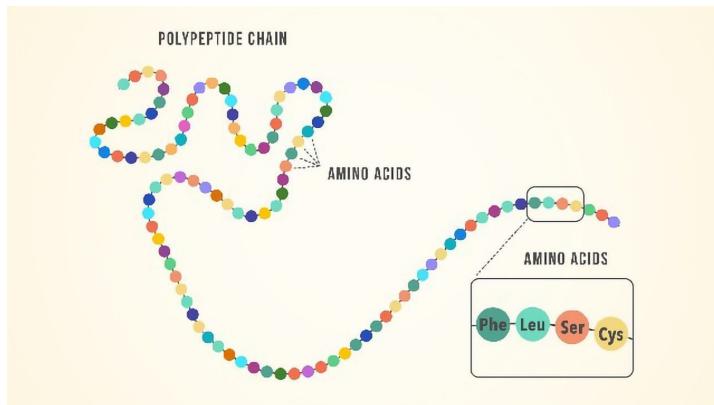
Polypeptide Chain

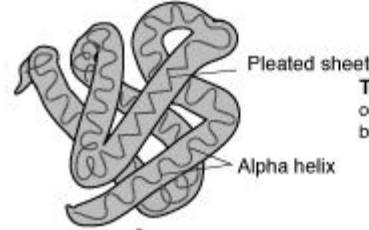
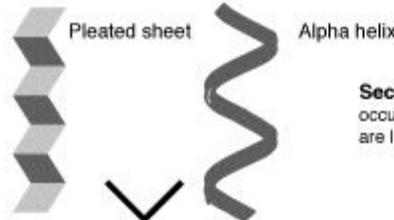
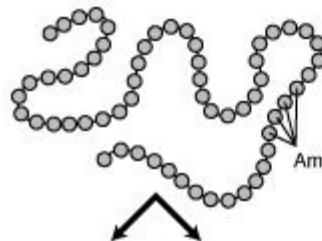
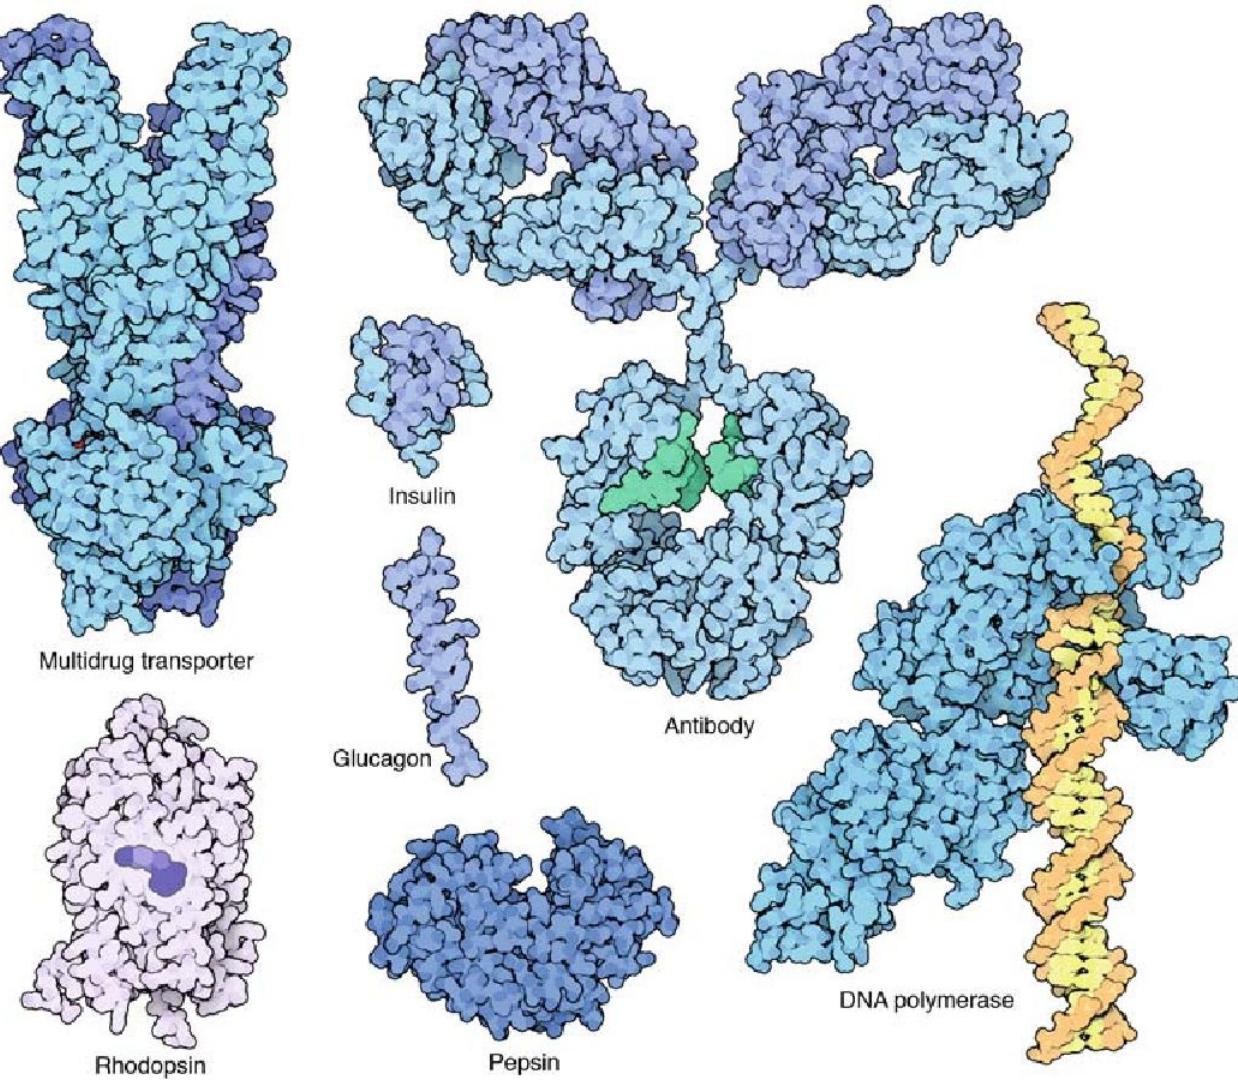


Protein folding problem

Protein's amino acid sequence -> three-dimensional atomic structure prediction.

The notion of a folding “problem” first emerged around 1960, with the appearance of the first atomic-resolution protein structures



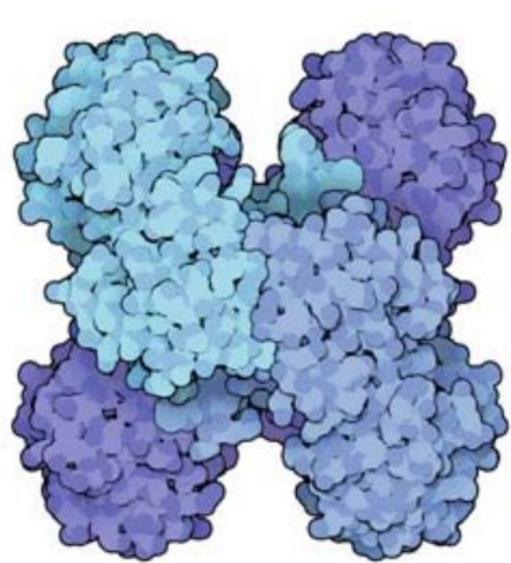


Primary protein structure
is sequence of a chain of amino acids

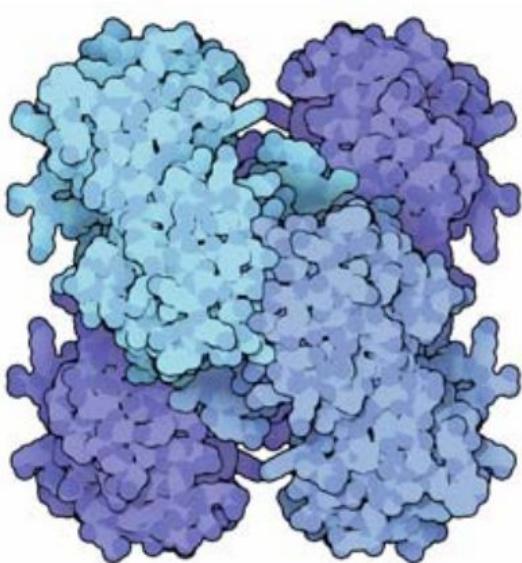
Secondary protein structure
occurs when the sequence of amino acids are linked by hydrogen bonds

Tertiary protein structure
occurs when certain attractions are made between alpha helices and pleated sheets

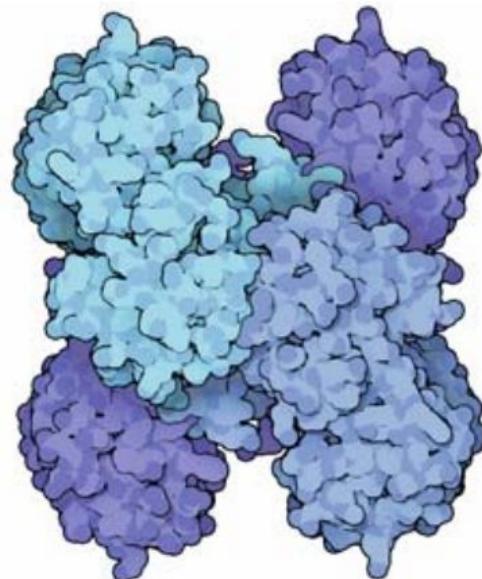
Quaternary protein structure
is a protein consisting of more than one amino acid chain.



Escherichia coli bacteria



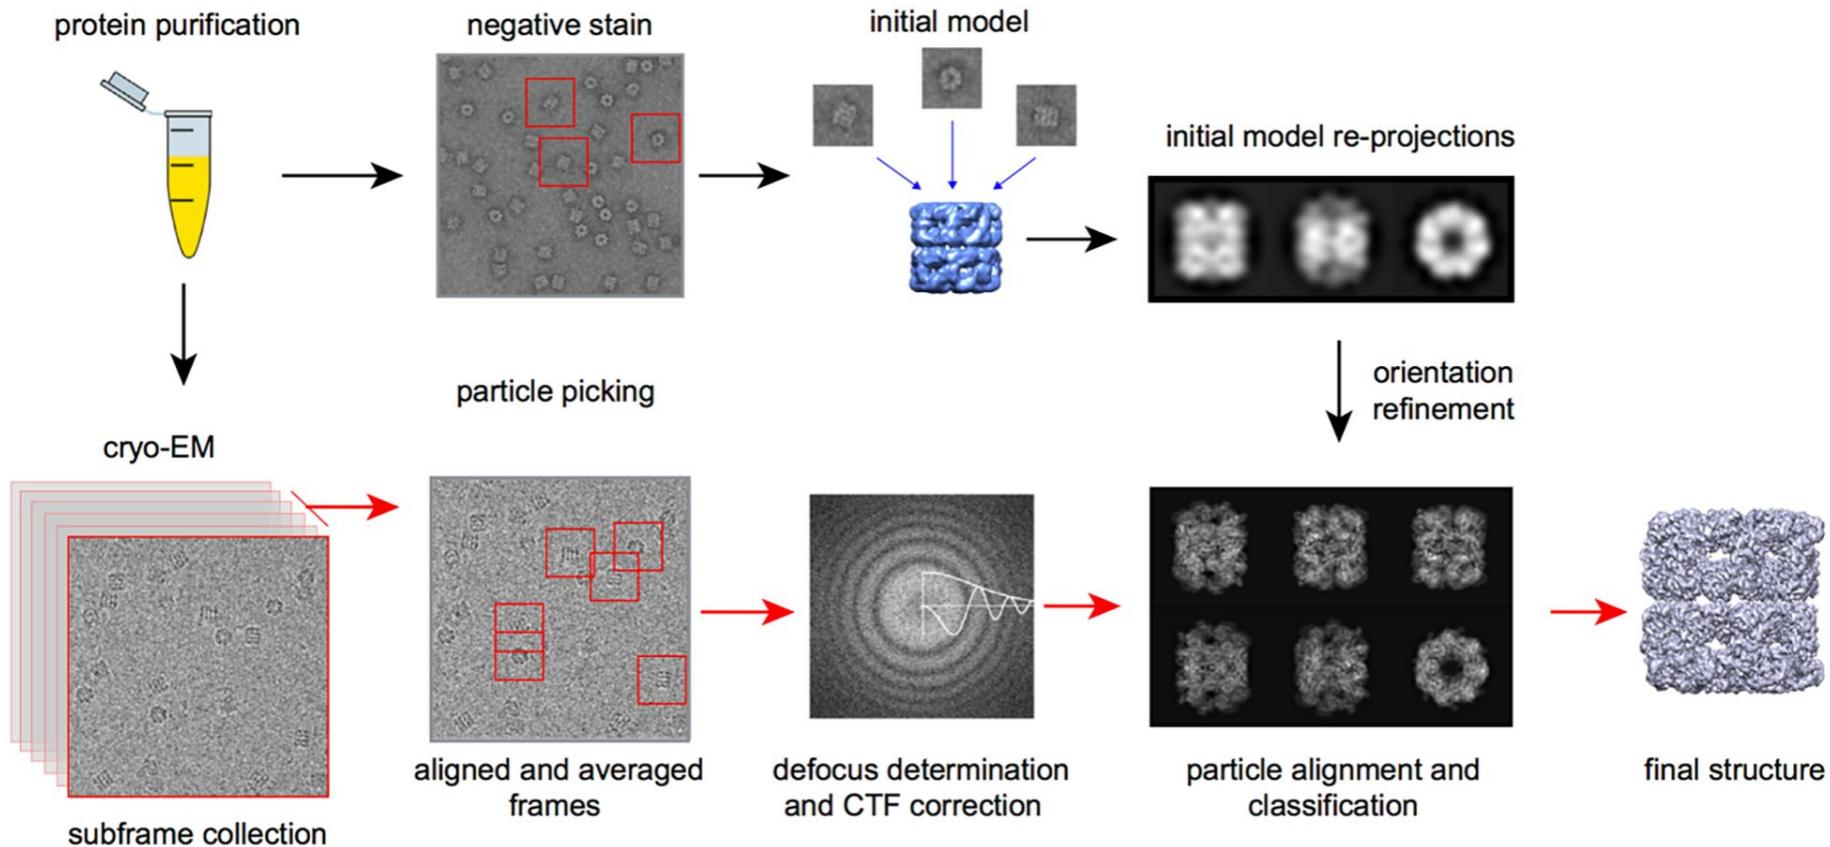
Spinach



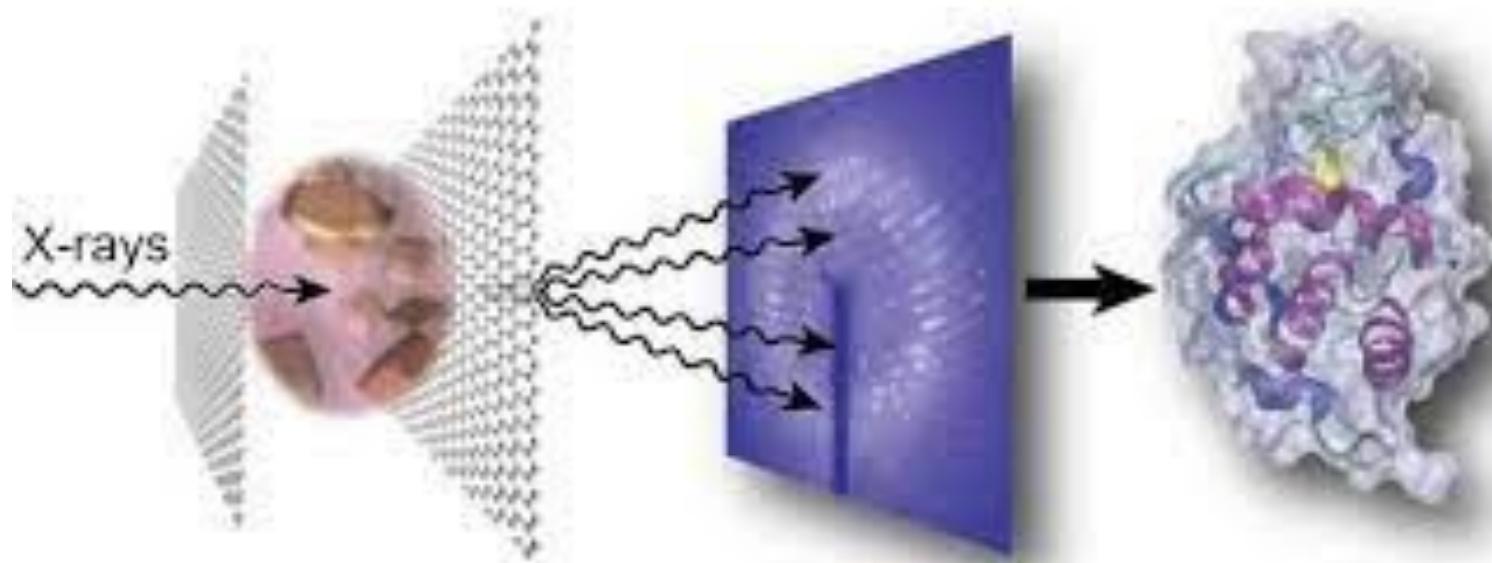
Human

Fig. 1.2 Molecular Machinery Many molecular machines are virtually identical in all living cells. This is particularly true for molecules that play an essential role in the processes of life, such as the enzyme glyceraldehyde-3-phosphate dehydrogenase, which is vital for the metabolism of sugar in all three organisms. This illustration shows the similar form of the enzyme from a bacterial cell (left), a plant cell (center), and human cells (right) (5,000,000 X)

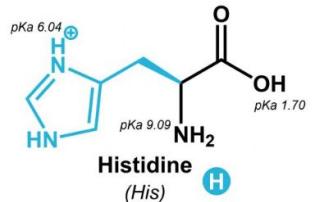
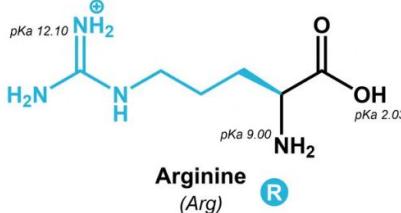
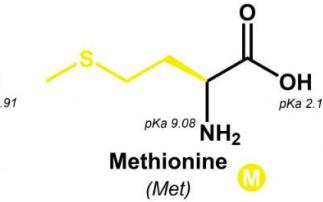
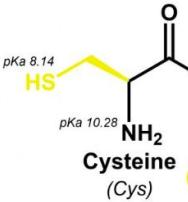
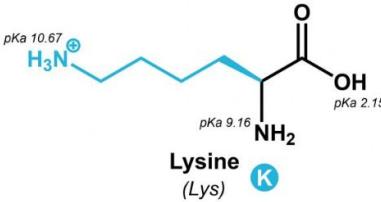
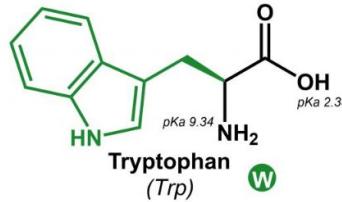
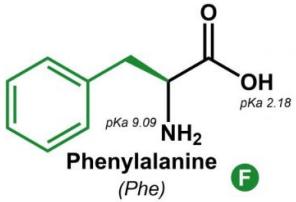
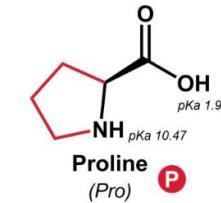
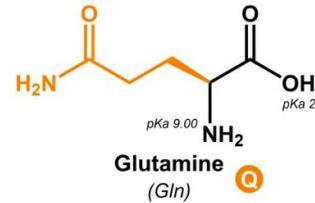
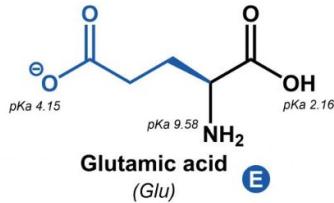
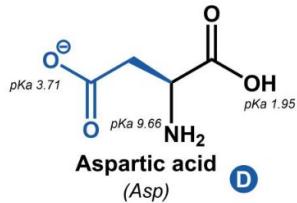
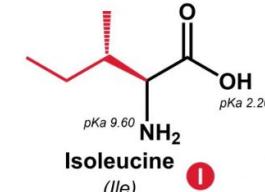
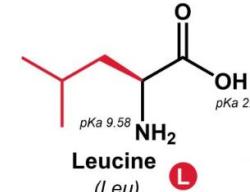
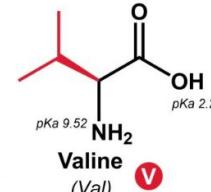
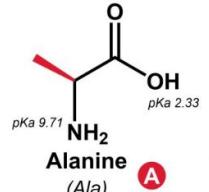
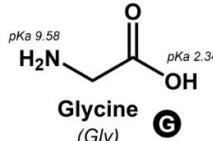
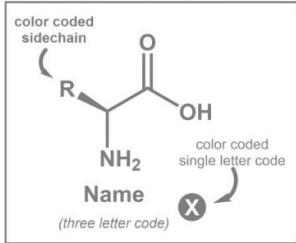
How to get Protein Structures: Cryo-EM



How to get Protein Structures: Crystallography



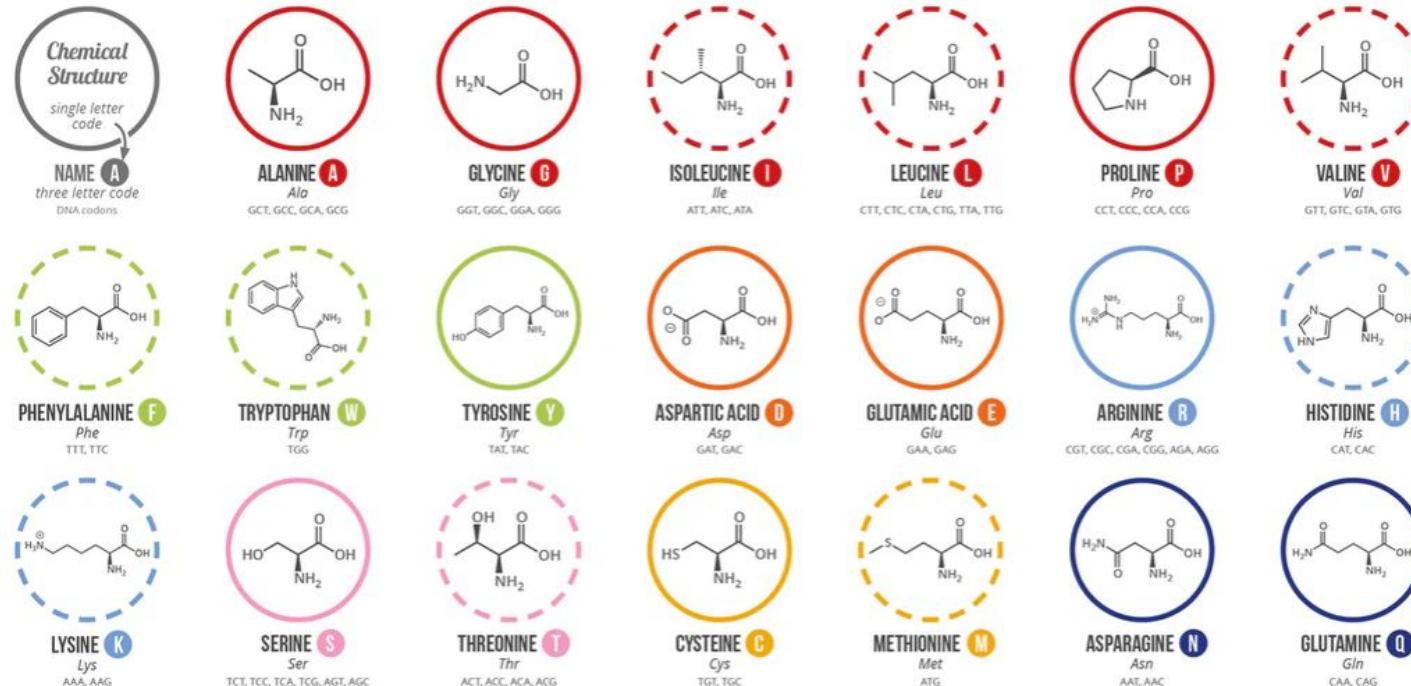
THE 20 COMMON AMINO ACIDS



A GUIDE TO THE TWENTY COMMON AMINO ACIDS

AMINO ACIDS ARE THE BUILDING BLOCKS OF PROTEINS IN LIVING ORGANISMS. THERE ARE OVER 500 AMINO ACIDS FOUND IN NATURE - HOWEVER, THE HUMAN GENETIC CODE ONLY DIRECTLY ENCODES 20. 'ESSENTIAL' AMINO ACIDS MUST BE OBTAINED FROM THE DIET, WHILST NON-ESSENTIAL AMINO ACIDS CAN BE SYNTHESISED IN THE BODY.

Chart Key:  ALIPHATIC  AROMATIC  ACIDIC  BASIC  HYDROXYLIC  SULFUR-CONTAINING  AMIDIC  NON-ESSENTIAL  ESSENTIAL



Note: This chart only shows those amino acids for which the human genetic code directly codes for. Selenocysteine is often referred to as the 21st amino acid, but is encoded in a special manner. In some cases, distinguishing between asparagine/aspartic acid and glutamine/glutamic acid is difficult. In these cases, the codes asx (B) and glx (Z) are respectively used.



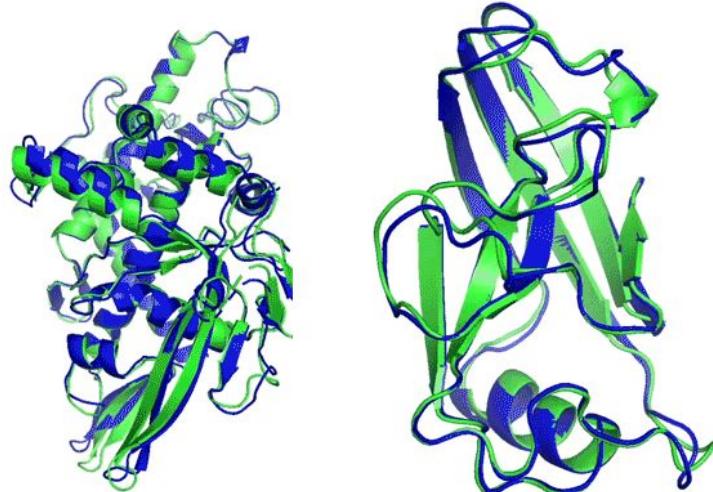
© COMPOUND INTEREST 2014 - WWW.COMPOUNDCHEM.COM | Twitter: @compoundchem | Facebook: www.facebook.com/compoundchem

Shared under a Creative Commons Attribution-NonCommercial-NoDerivatives licence.



Critical Assessment of protein Structure Prediction (CASP)

- Biennial **global blind-prediction experiment** (since 1994)
- Fully blind evaluation → no data leakage → highest credibility



Protein folding example, where predictions (in blue) are aligned to the known protein structure (in green) in the Protein DataBase (PDB). (image source: DeepMind blog)

David Baker's Lab

- Professor at Univ Washington
- Known for **computational Protein Design**
- **Rosetta** (physics-based)
- **RoseTTAFold** (ML for protein structure prediction)
- **Protein MPNN** (ML for protein inverse folding)
- **RF diffusion** (generative model for protein)

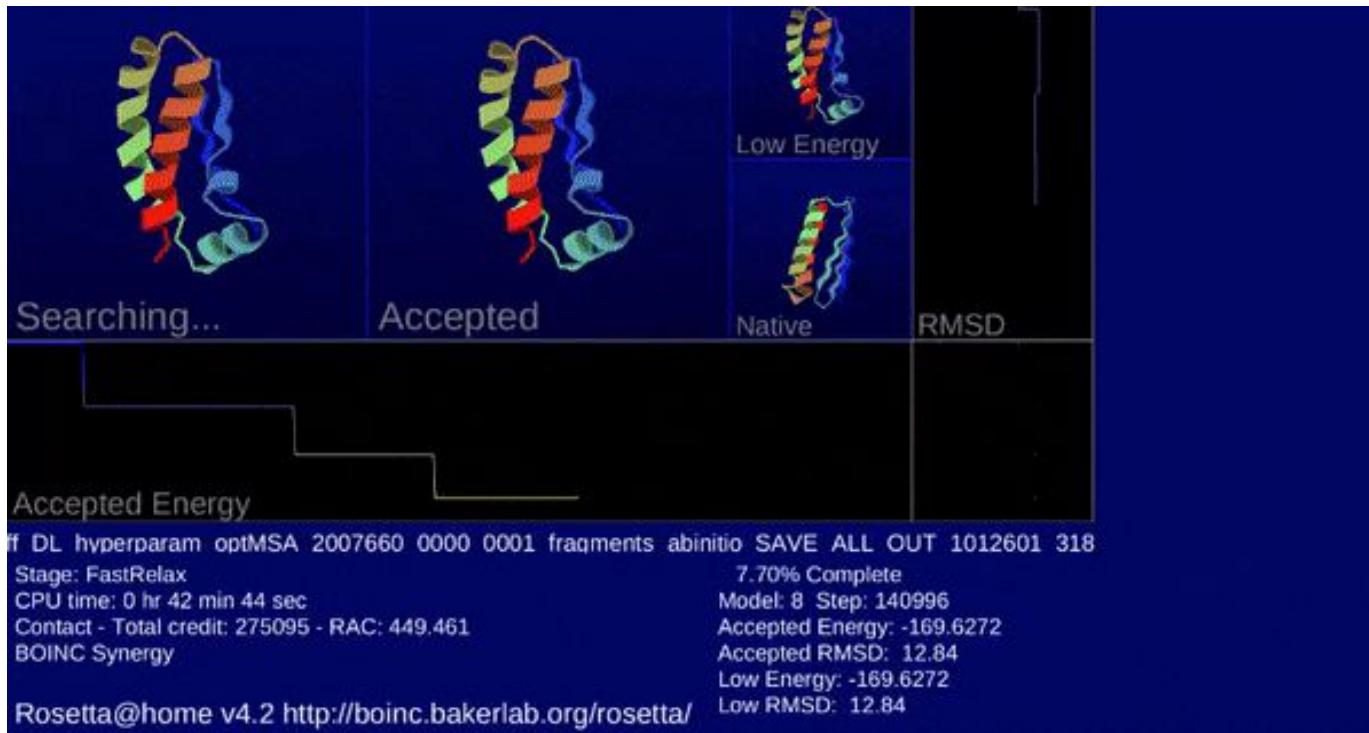


David Baker



David Baker's Lab

Rosetta (Physics based structure prediction)



Andrew Leaver-Fay
& Rosetta Com

AlphaFold II



John Jumper (DeepMind)

Article

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

 Check for updates

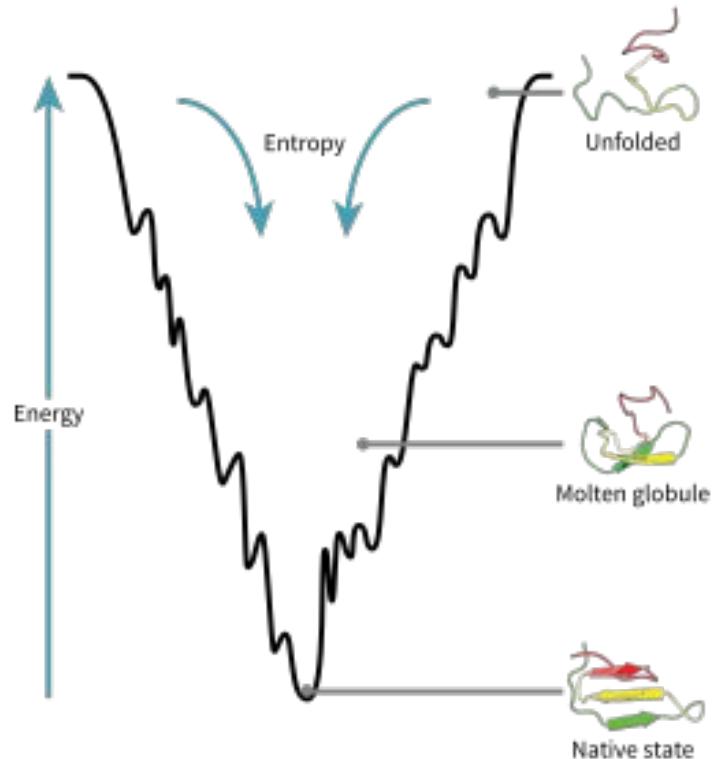
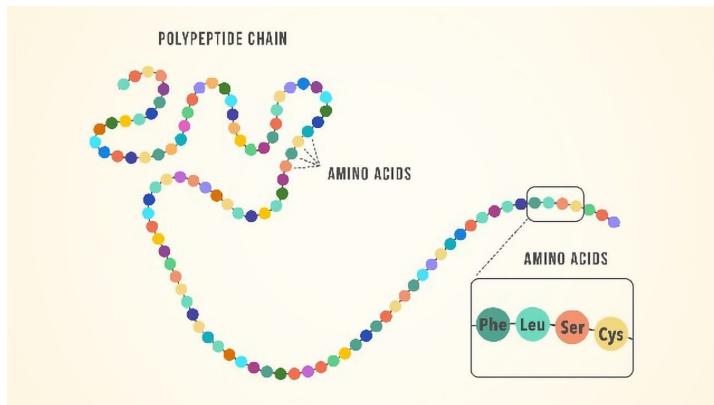
John Jumper^{1,4} , Richard Evans^{1,4}, Alexander Pritzel^{1,4}, Tim Green^{1,4}, Michael Figurnov^{1,4}, Olaf Ronneberger^{1,4}, Kathryn Tunyasuvunakool^{1,4}, Russ Bates^{1,4}, Augustin Žídek^{1,4}, Anna Potapenko^{1,4}, Alex Bridgland^{1,4}, Clemens Meyer^{1,4}, Simon A. A. Kohl^{1,4}, Andrew J. Ballard^{1,4}, Andrew Cowie^{1,4}, Bernardino Romera-Paredes^{1,4}, Stanislav Nikolov^{1,4}, Rishabh Jain^{1,4}, Jonas Adler¹, Trevor Back¹, Stig Petersen¹, David Reiman¹, Ellen Clancy¹, Michal Zielinski¹, Martin Steinegger^{2,3}, Michalina Pacholska¹, Tamas Berghammer¹, 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.

Protein folding problem

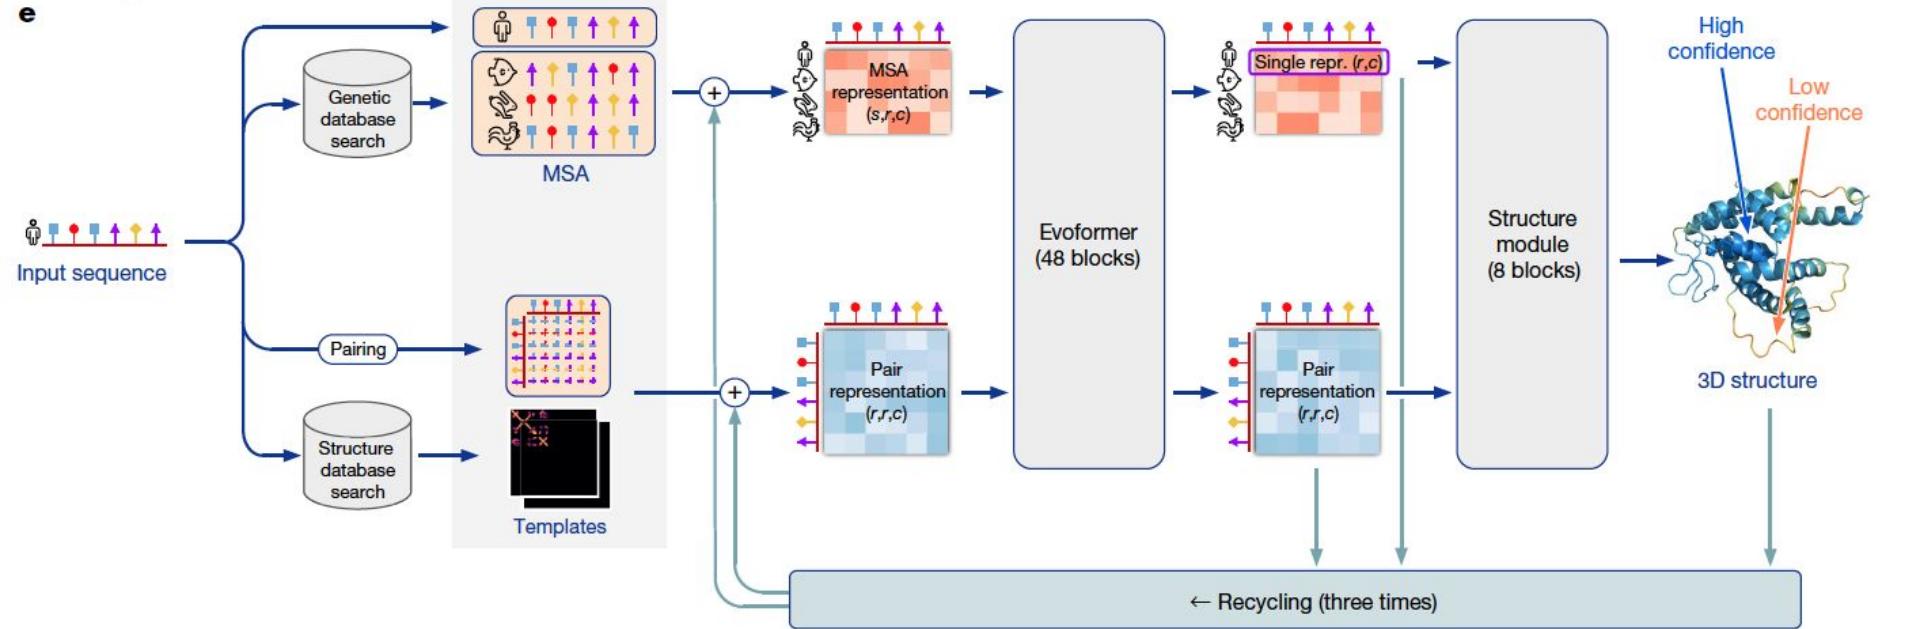
Protein's amino acid sequence -> three-dimensional atomic structure prediction.

The notion of a folding “problem” first emerged around 1960, with the appearance of the first atomic-resolution protein structures



Network

AlphaFold2

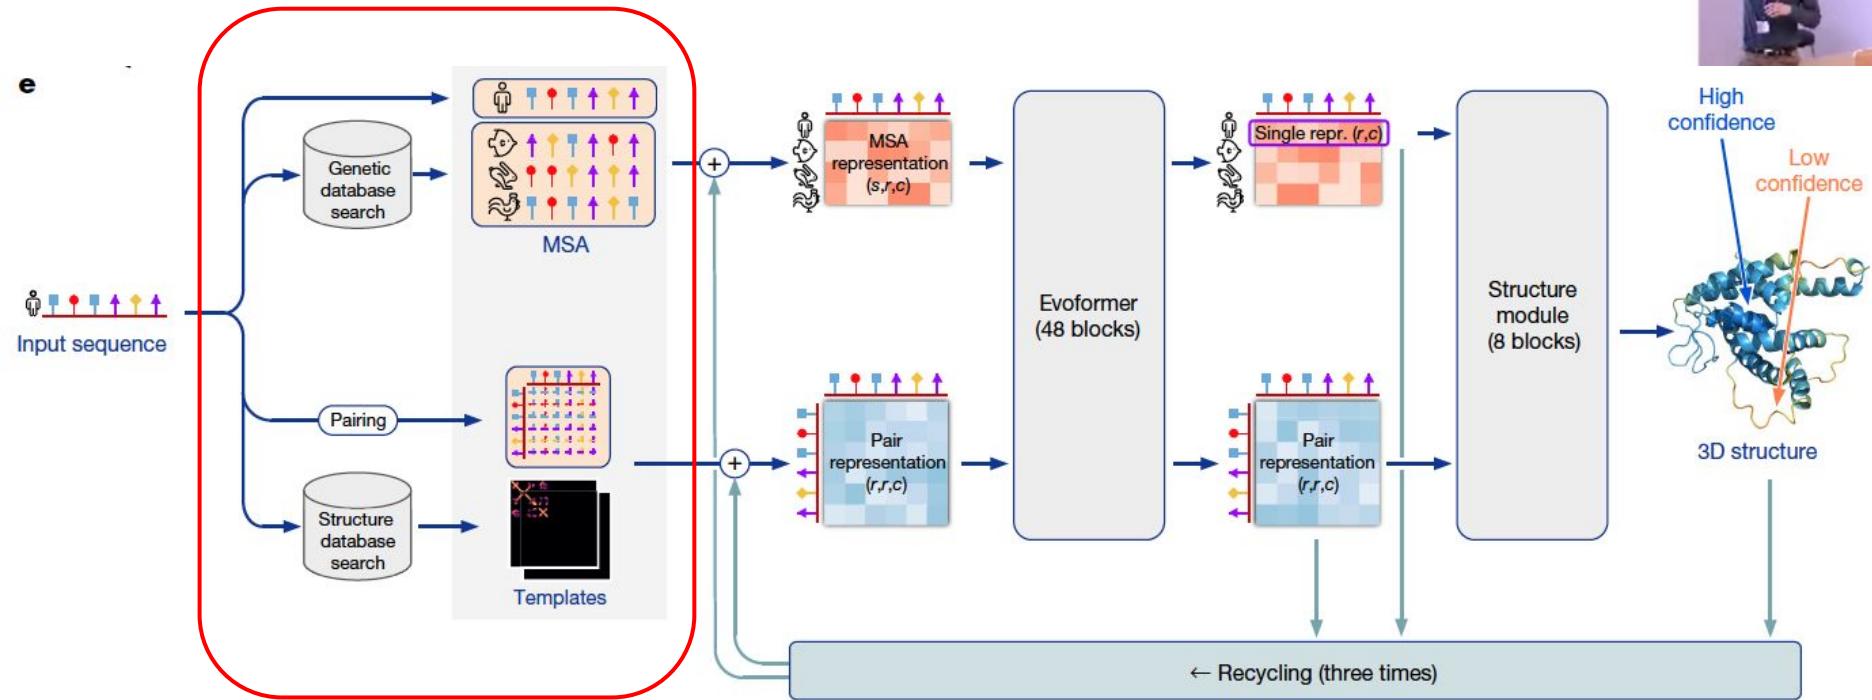


The Structure Module predicts a rotation + translation to place each residue.

A small network predicts side chain chi angles. The final structure is run through a relaxation process.



Network AlphaFold2

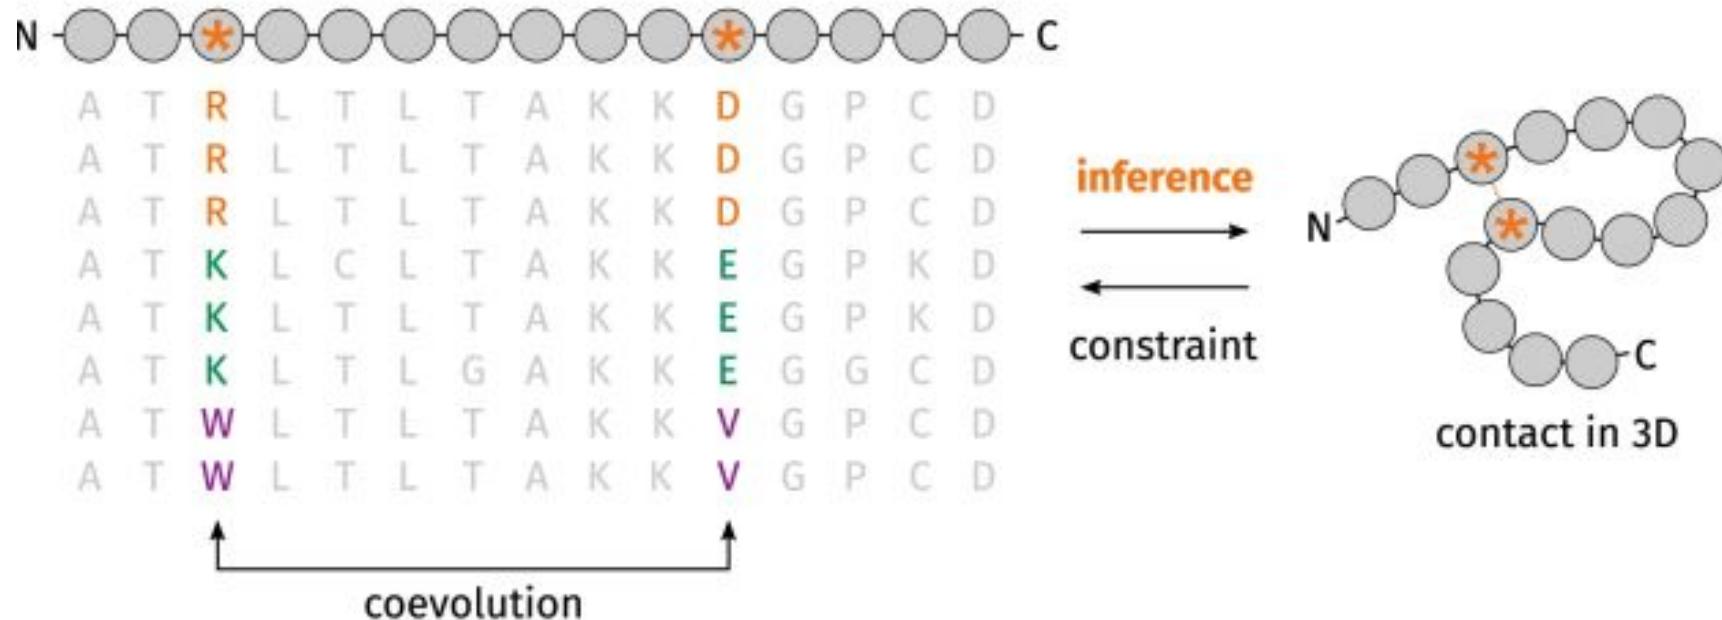


The Structure Module predicts a rotation + translation to place each residue.

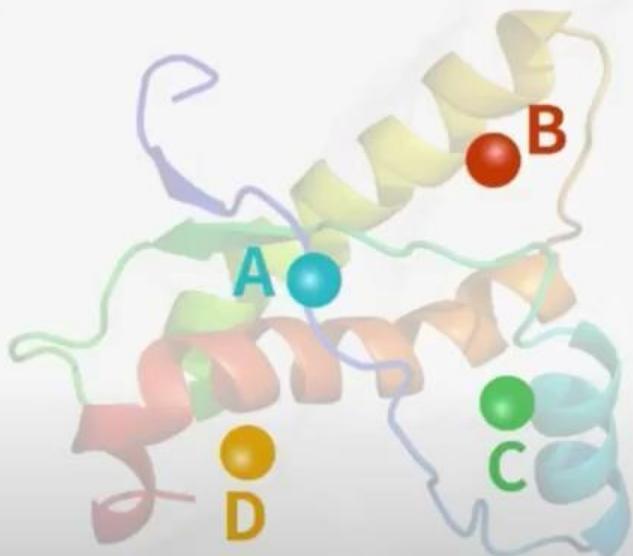
A small network predicts side chain chi angles. The final structure is run through a relaxation process.



Multiple Sequence Alignment (MSA)



Distance Matrix



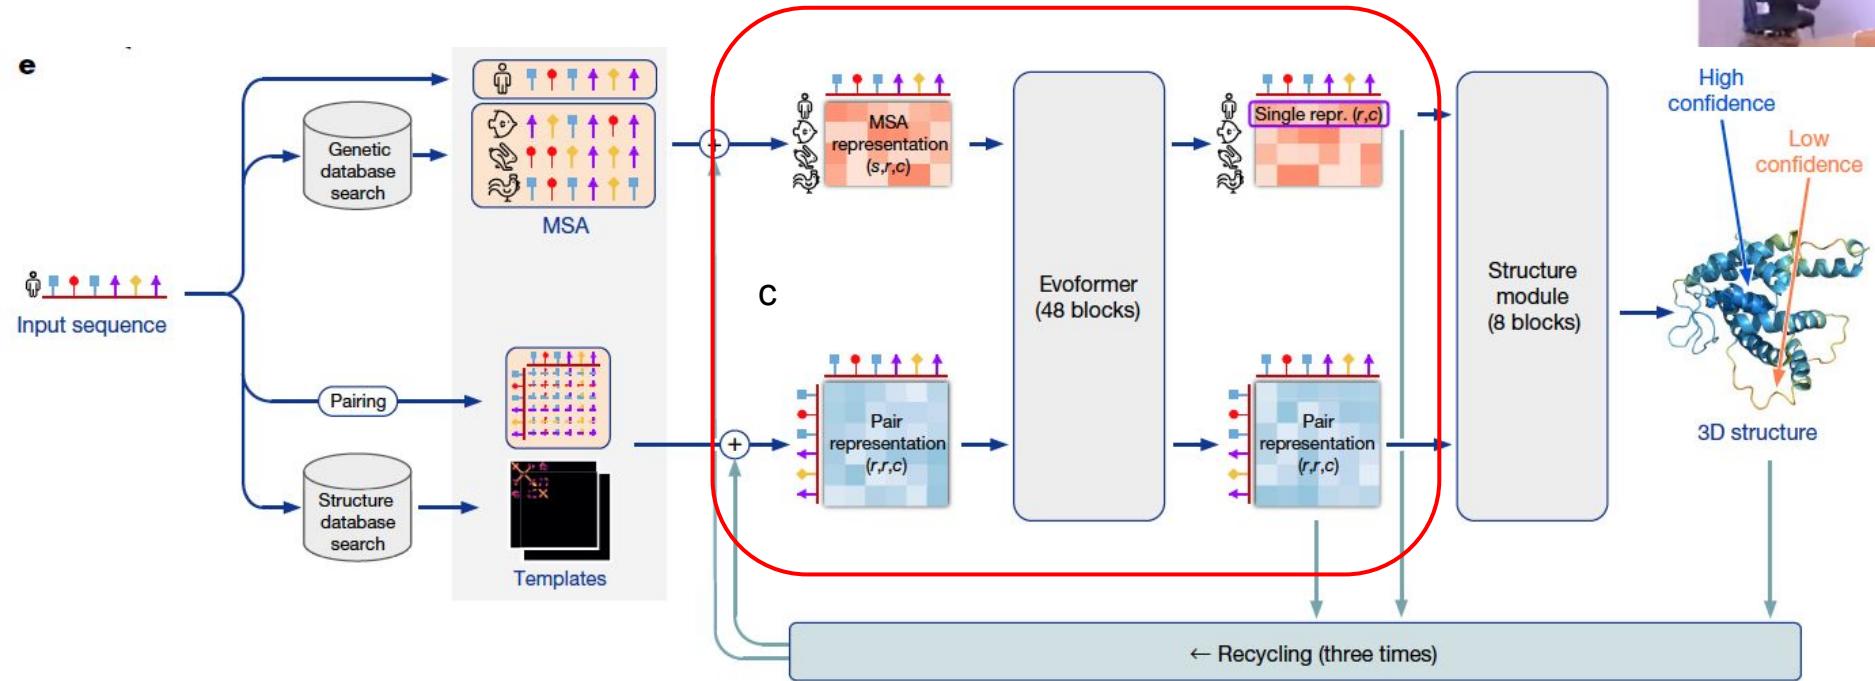
	A	B	C	D
A	0			
B		0		
C			0	
D				0

Network

AlphaFold2



e

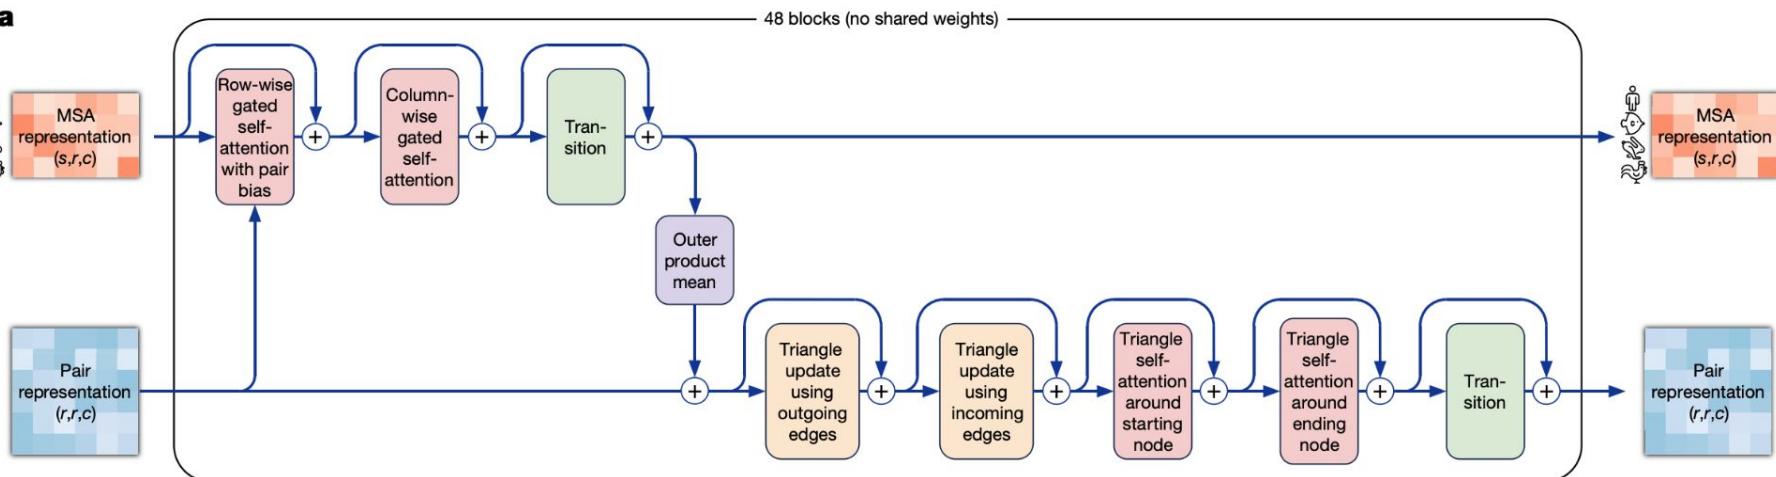


The Structure Module predicts a rotation + translation to place each residue.

A small network predicts side chain chi angles. The final structure is run through a relaxation process.

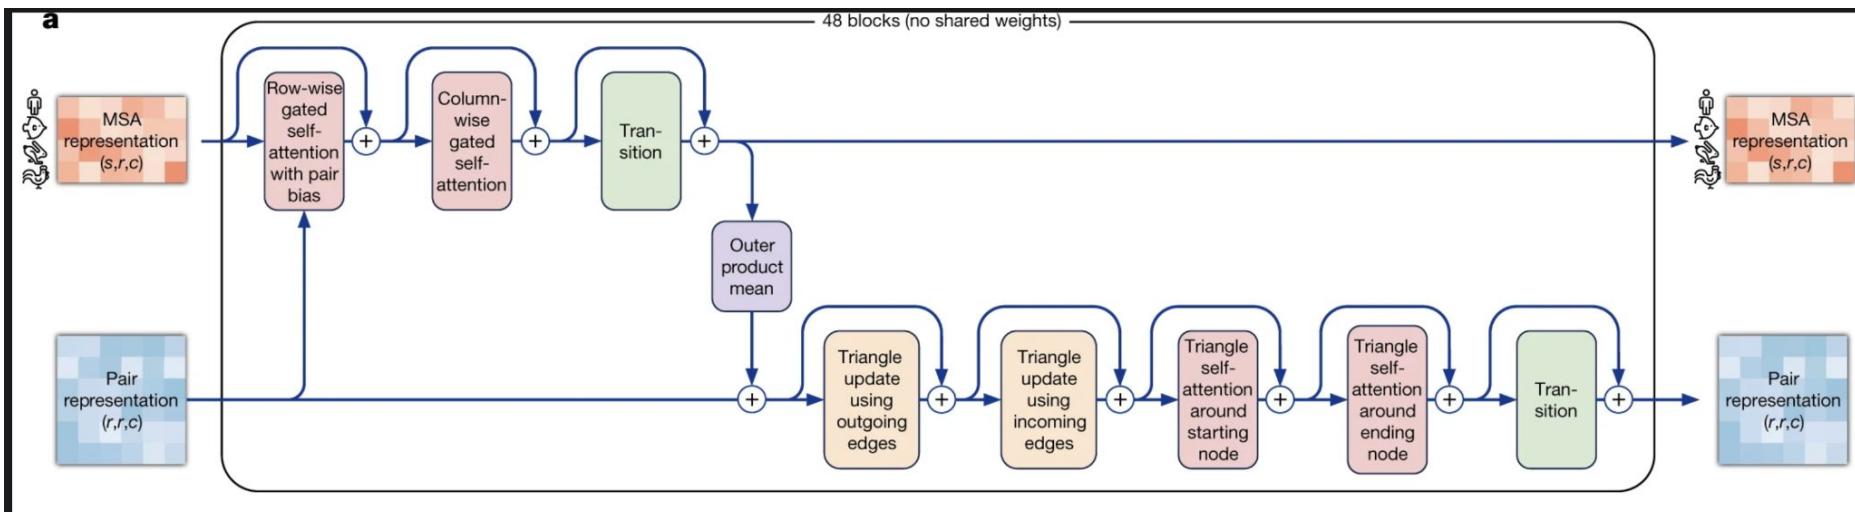
Evoformer

a

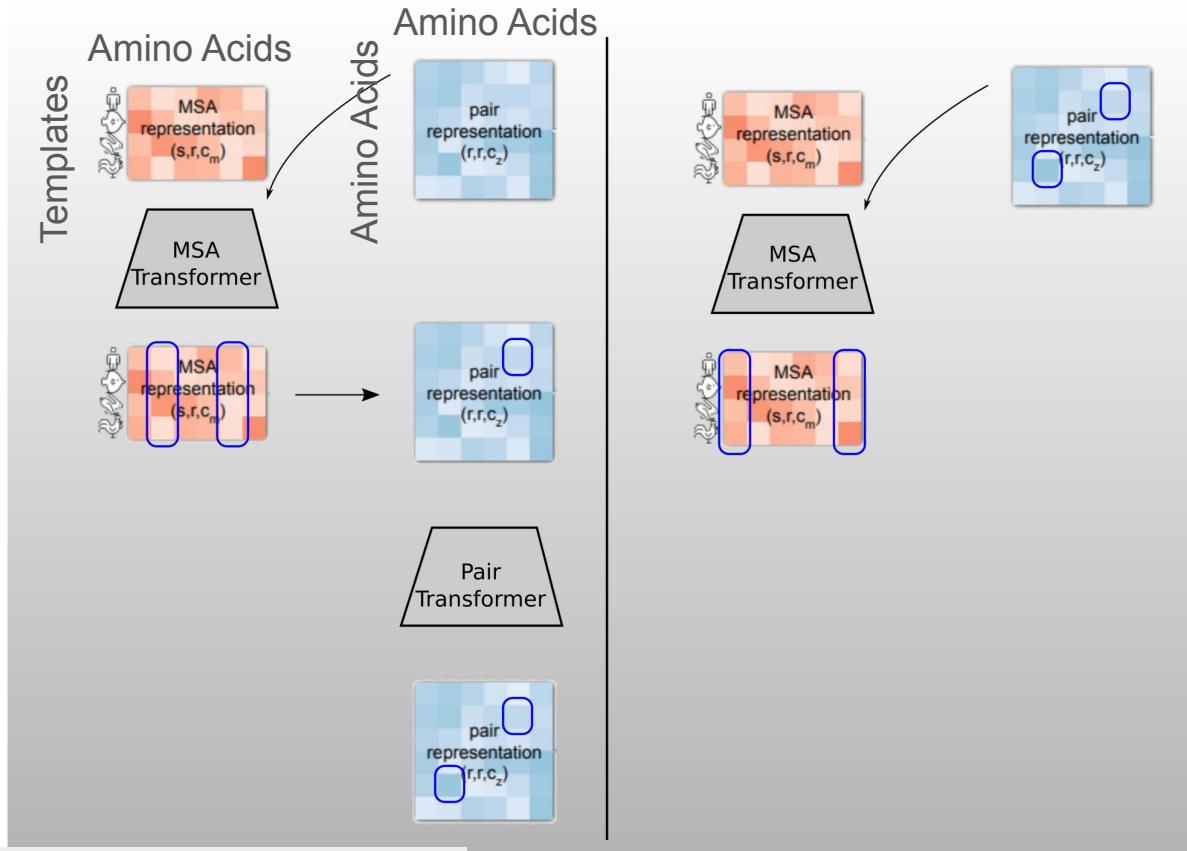


What is going on in the swap?

1D of many proteins and 2D of one protein -> Update both

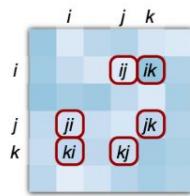


Representations: Swap between 1D and 2D

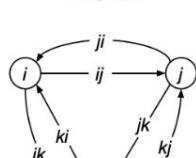


Triangular Attention

b Pair representation
(r, r, c)

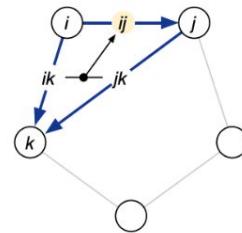


Corresponding edges
in a graph

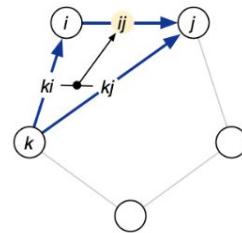


c

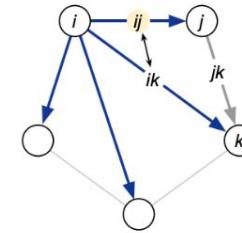
Triangle multiplicative update
using 'outgoing' edges



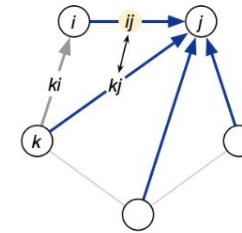
Triangle multiplicative update
using 'incoming' edges



Triangle self-attention around
starting node



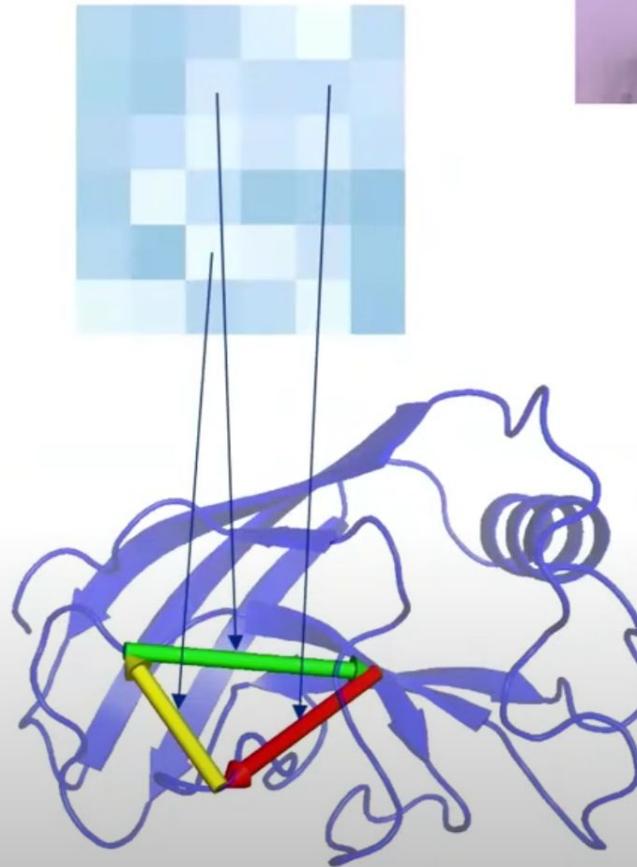
Triangle self-attention around
ending node



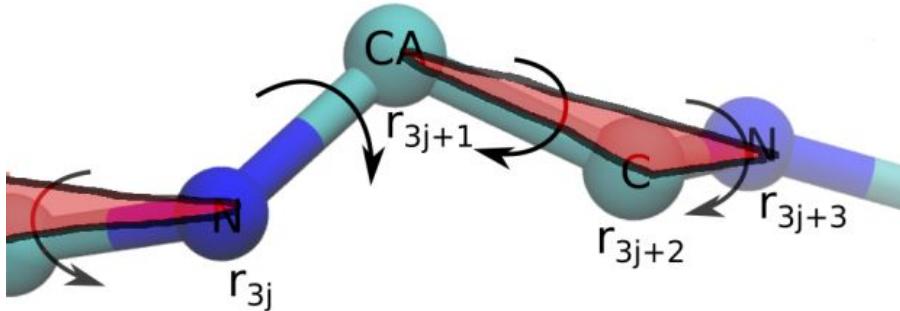
Vignette 1: Triangular Attention



- **Take 3 points A, B, C**
 - ◆ If Distance AB and distance BC known strong constraint on AC (triangle inequality)
 - ◆ Evolution & Sequence gives information about relations between residues
- **Pair Embedding encodes relations**
 - ◆ Update for pair AC should depend on BC, AB
 - ◆ All about who communicates in the network, not what is computed



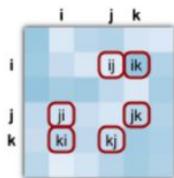
What is going on in the swap?



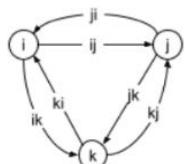
Concepts: Affine Transforms, Triangle Inequality

We preserve triangle inequality but *don't* care about overlaps till the very end.

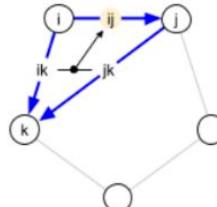
b pair representation
(r, r, c')



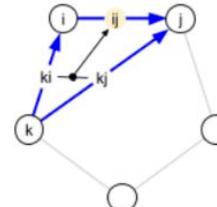
corresponding edges
in a graph



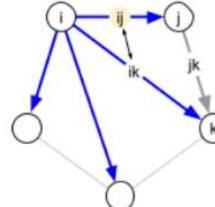
c Triangle multiplicative update
using "outgoing" edges



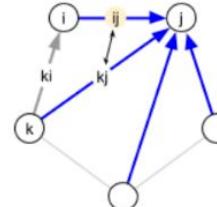
Triangle multiplicative update
using "incoming" edges



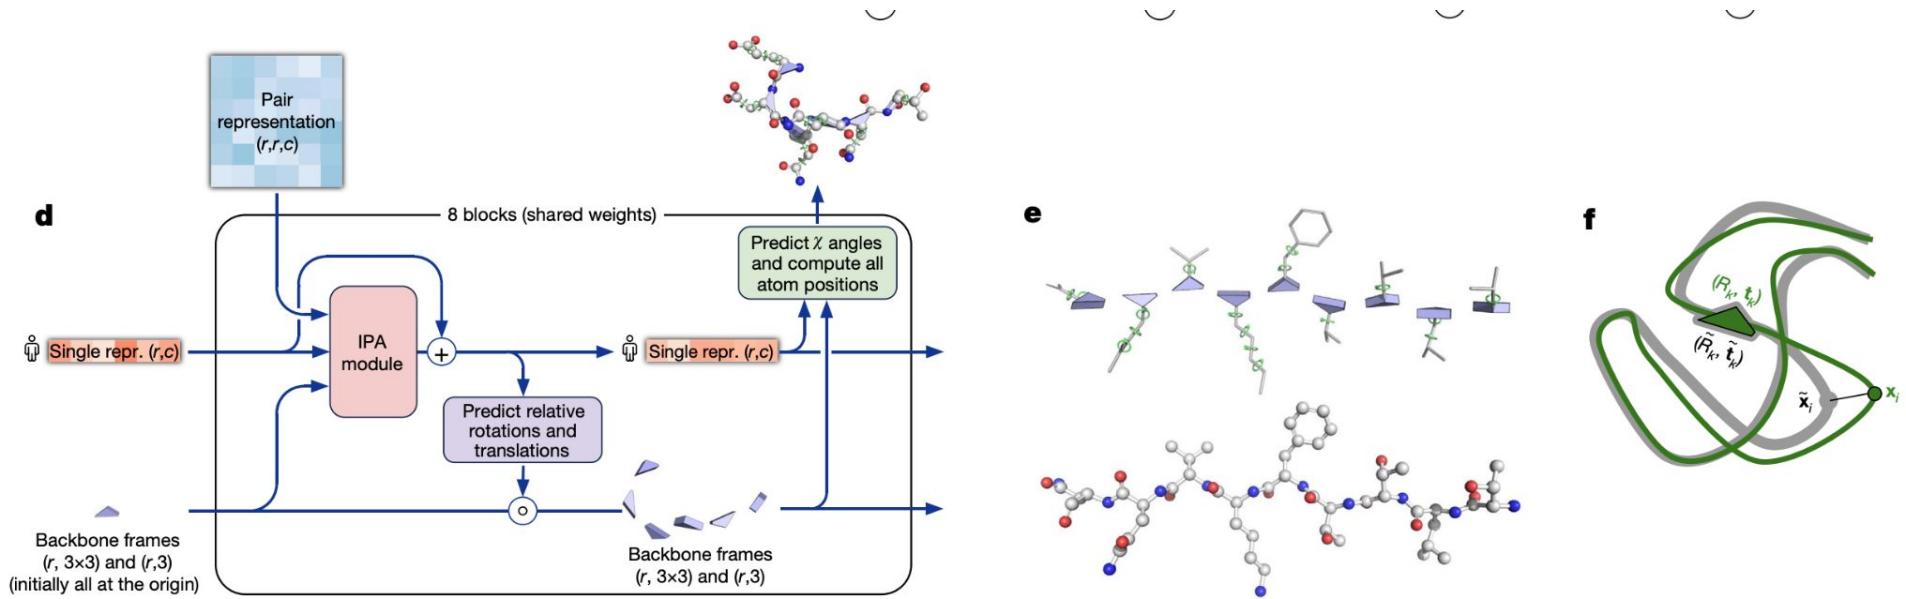
Triangle self-attention around
starting node



Triangle self-attention around
ending node

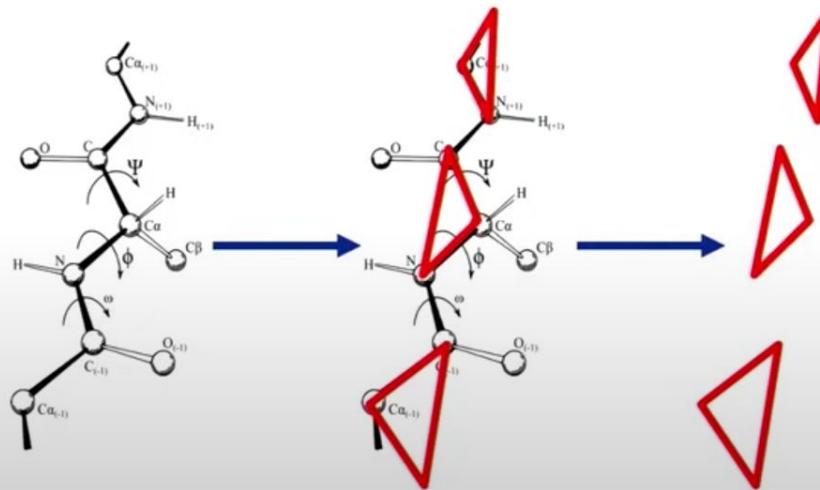


Structure Prediction



Vignette 2: Structure module

- **End-to-end folding** instead of gradient descent
- Protein backbone = gas of 3-D rigid bodies
(chain is learned!)



- **3-D equivariant transformer architecture** updates the rigid bodies / backbone
 - Also builds the side chains from torsion angles

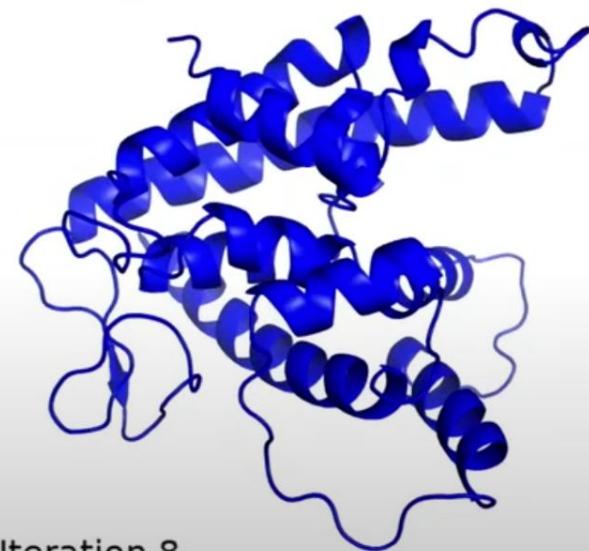
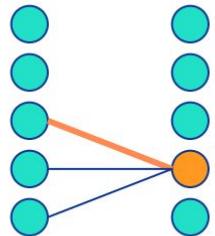


Image: Dcrjsr, vectorised Adam Rędzikowski (CC BY 3.0, Wikipedia)

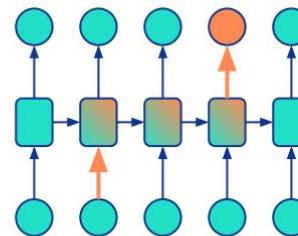
Inductive Bias for Deep Learning Models

© 2020 DeepMind Technologies Limited



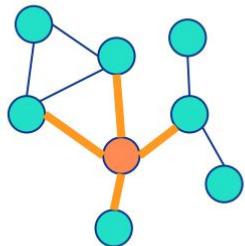
Convolutional Networks (e.g. computer vision)

- data in regular grid
- information flow to local neighbours



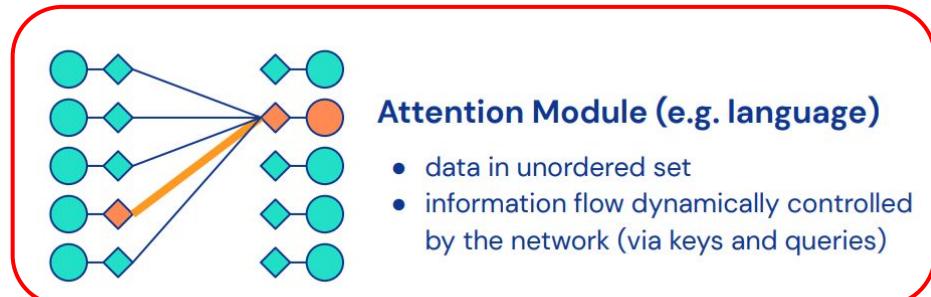
Recurrent Networks (e.g. language)

- data in ordered sequence
- information flow sequentially



Graph Networks (e.g. recommender systems or molecules)

- data in fixed graph structure
- information flow along fixed edges



Attention Module (e.g. language)

- data in unordered set
- information flow dynamically controlled by the network (via keys and queries)



Training loss functions

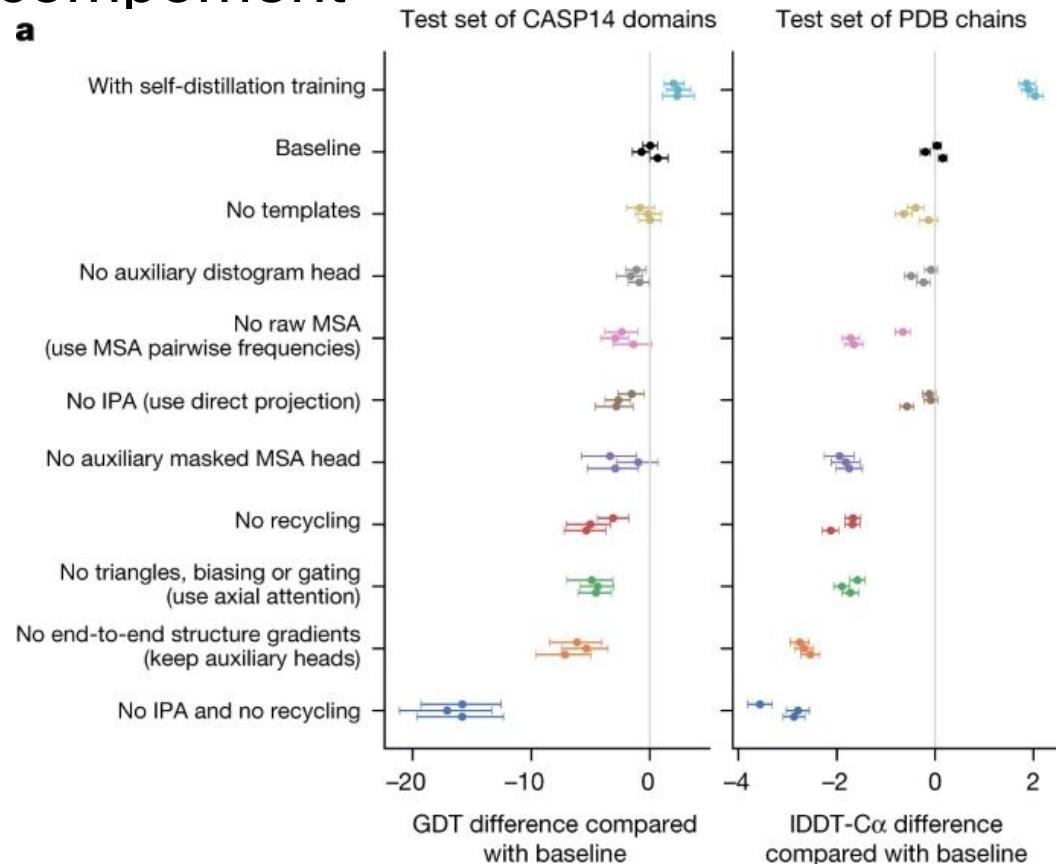
$$\mathcal{L} = \begin{cases} 0.5\mathcal{L}_{\text{FAPE}} + 0.5\mathcal{L}_{\text{aux}} + 0.3\mathcal{L}_{\text{dist}} + 2.0\mathcal{L}_{\text{msa}} + 0.01\mathcal{L}_{\text{conf}} & \text{training} \\ 0.5\mathcal{L}_{\text{FAPE}} + 0.5\mathcal{L}_{\text{aux}} + 0.3\mathcal{L}_{\text{dist}} + 2.0\mathcal{L}_{\text{msa}} + 0.01\mathcal{L}_{\text{conf}} + 0.01\mathcal{L}_{\text{exp resolved}} + 1.0\mathcal{L}_{\text{viol}} & \text{fine-tuning} \end{cases}$$

Training losses are designed to emphasize contributions from different architectural parts:

- \mathcal{L}_{aux} : averaged FAPE and torsion angle losses from intermediate structure module layers
- $\mathcal{L}_{\text{dist}}$: distogram prediction (cross-entropy)
- \mathcal{L}_{msa} : masked-MSA prediction (cross-entropy)

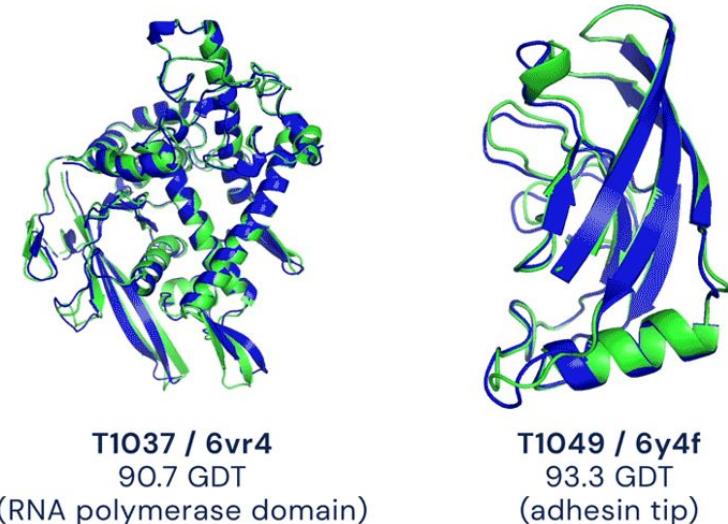
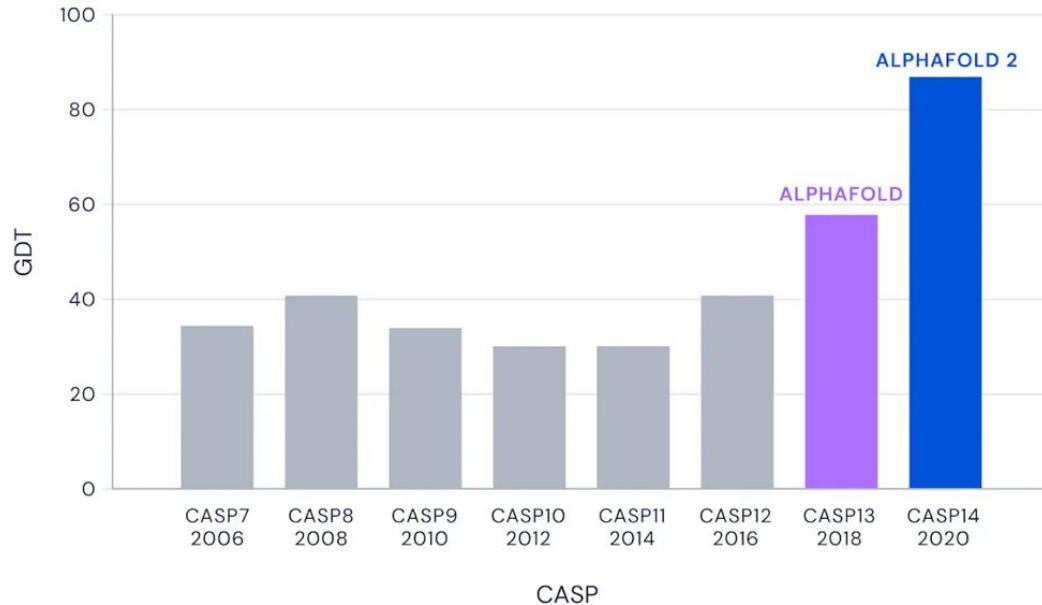
Ablation study of each component

a



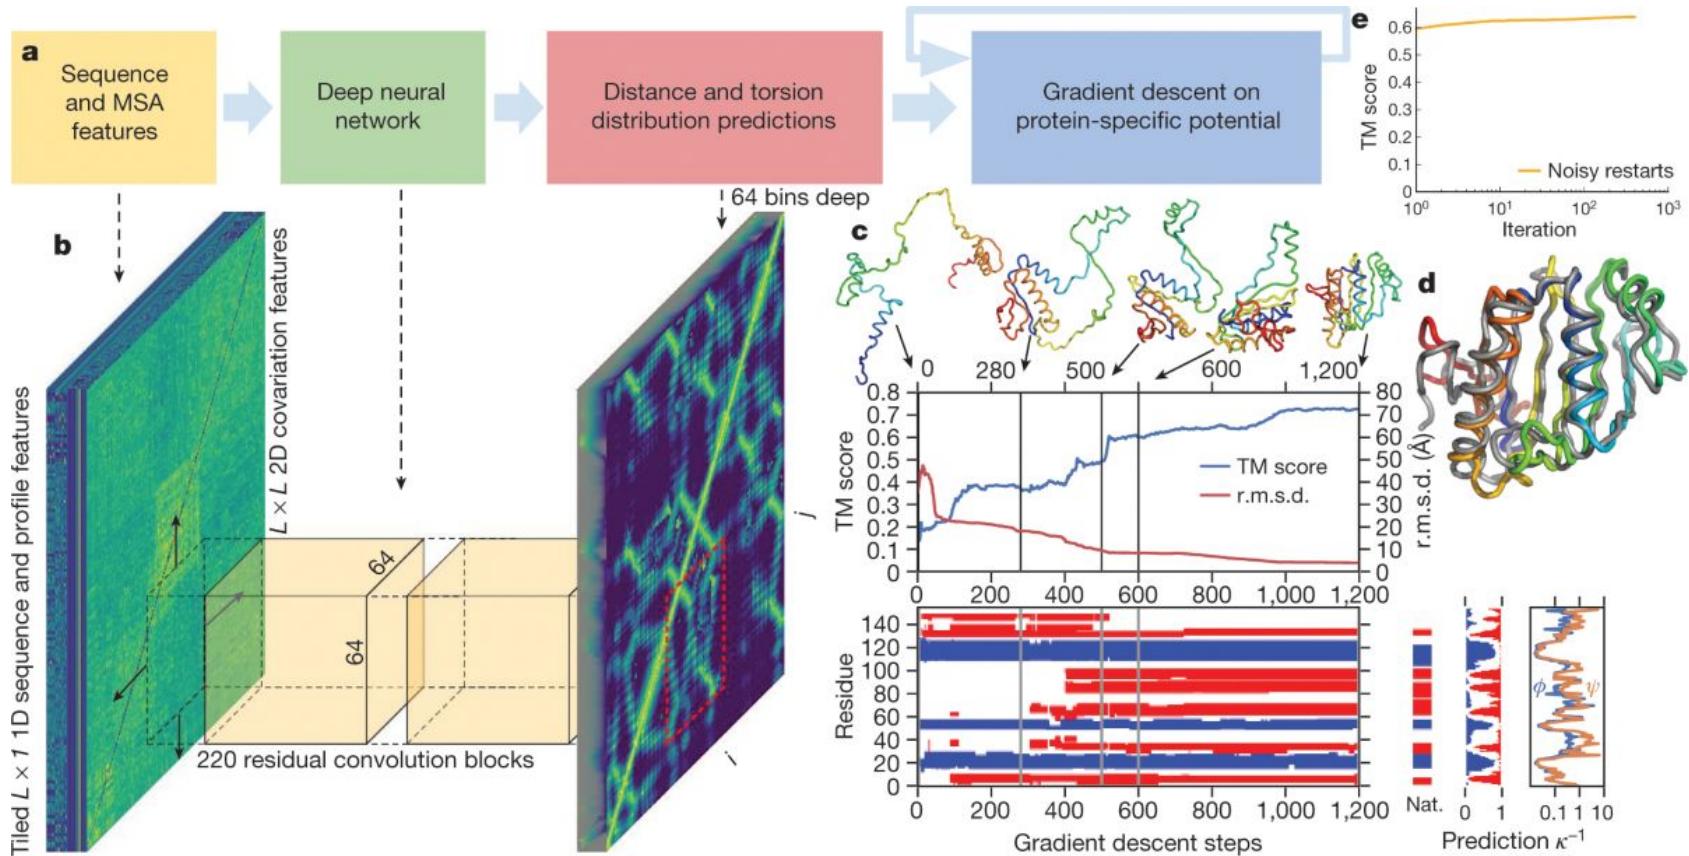
AlphaFold 2 in CASP challenge

Median Free-Modelling Accuracy



- Experimental result
- Computational prediction

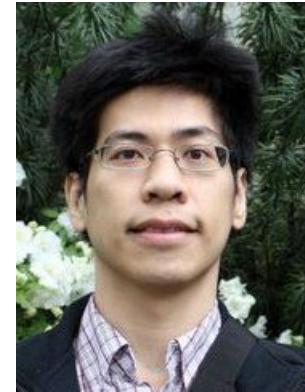
Comparison to AlphaFold 1



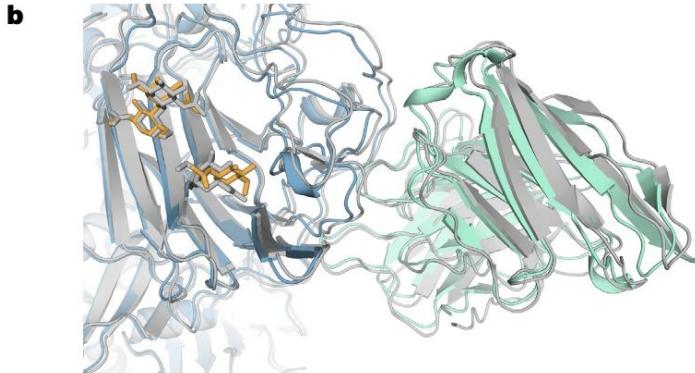
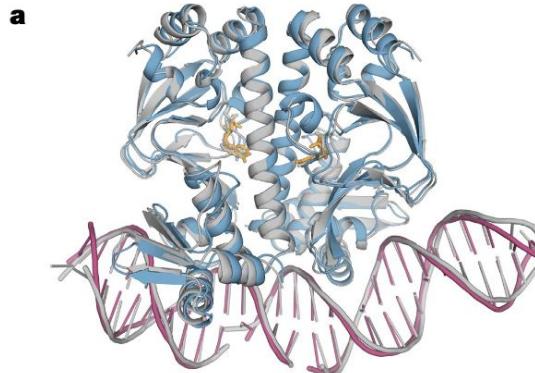
AlphaFold 2 vs AlphaFold 3 (improvement)

AF3 has these improvement compare with AF2

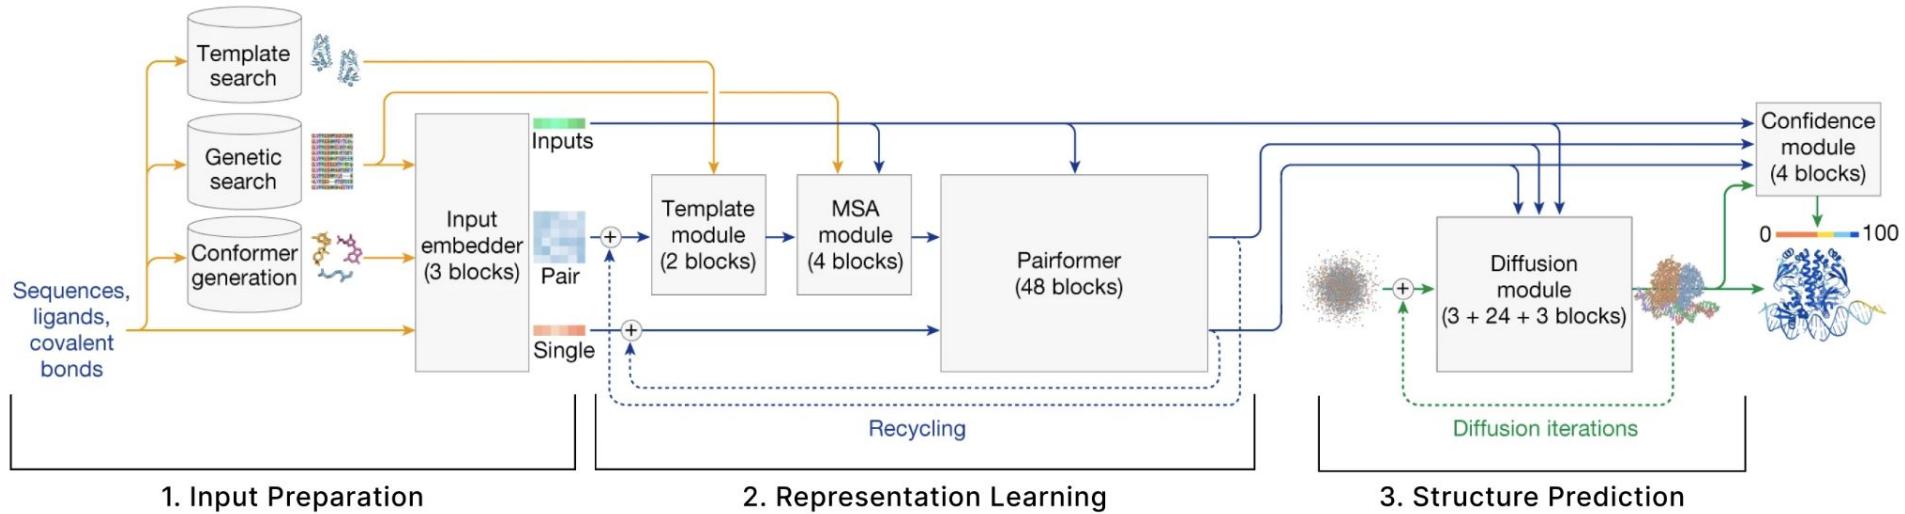
- multiple proteins
- DNA
- RNA
- small molecule ligands interaction with proteins
- Improved Antibody structures prediction
- Pseudo-opensource model
- Remove equivariant models, infuse diffusion model in protein structure prediction

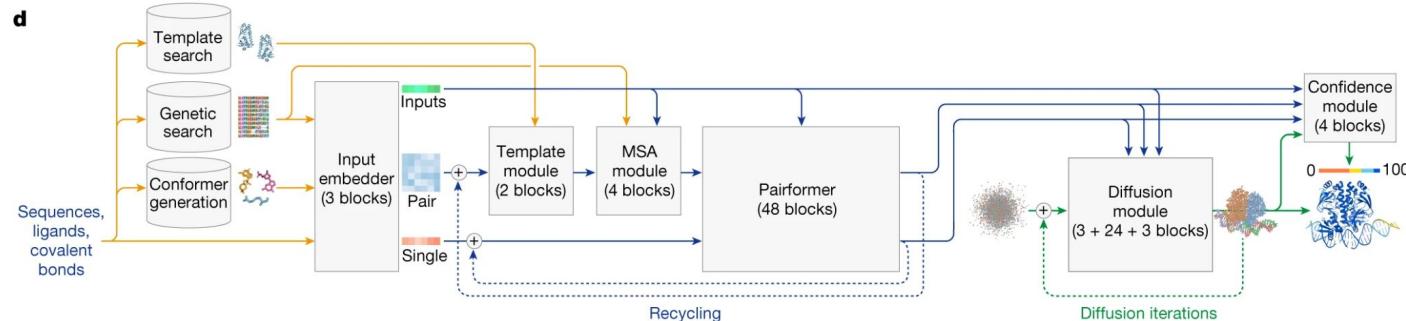
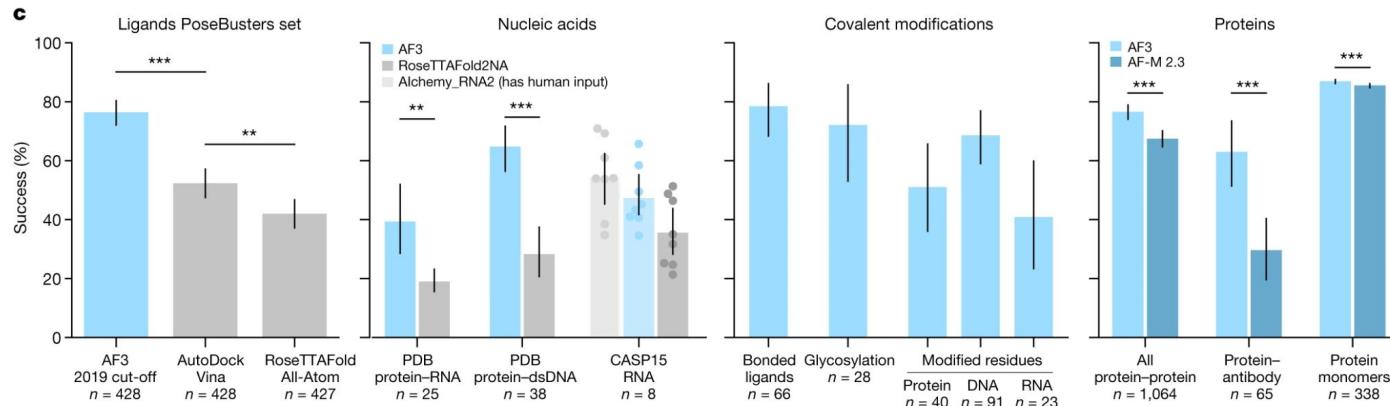
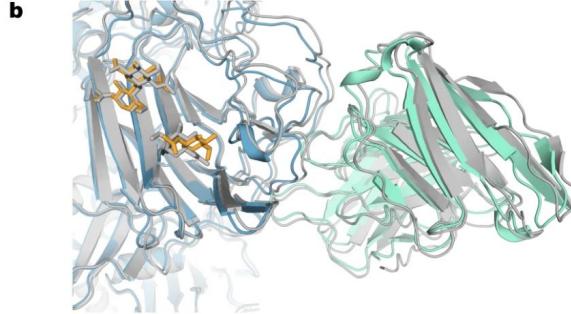
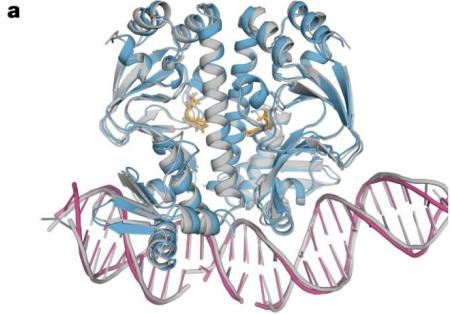


Chia-Chun (Alden) Hung
Isomorphic Lab



AlphaFold 3 architecture





Illustrations: Niklas Elmehed

THE NOBEL PRIZE IN CHEMISTRY 2024



David
Baker

"for computational
protein design"

Demis
Hassabis

"for protein structure prediction"

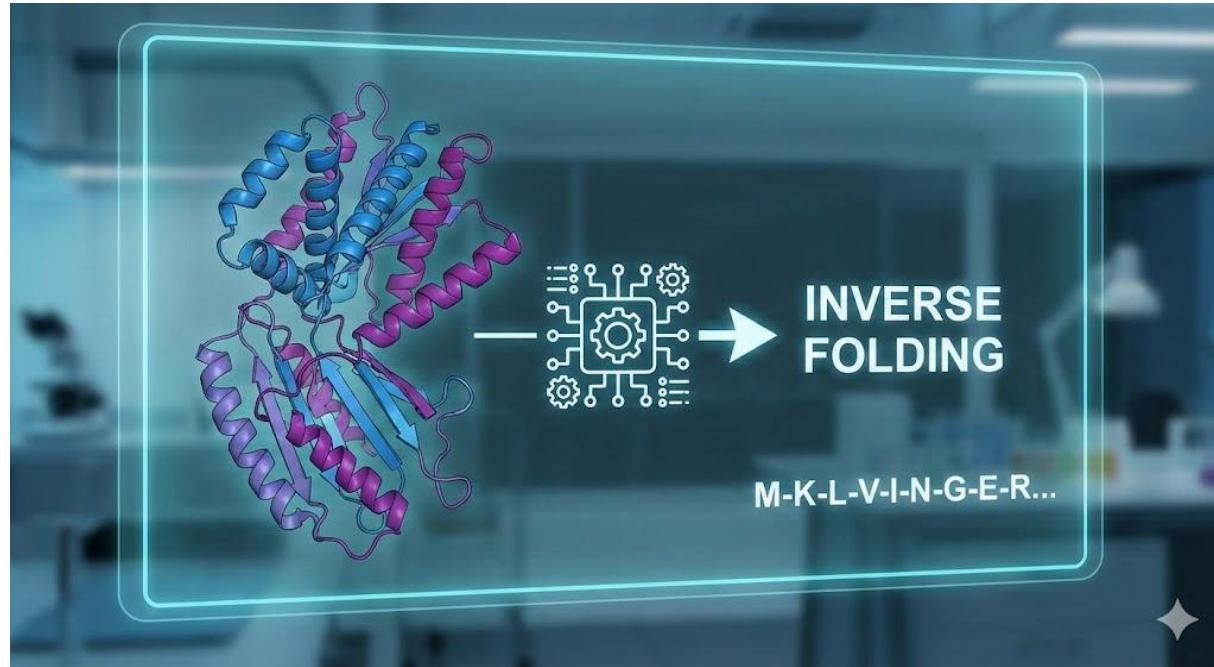
John M.
Jumper

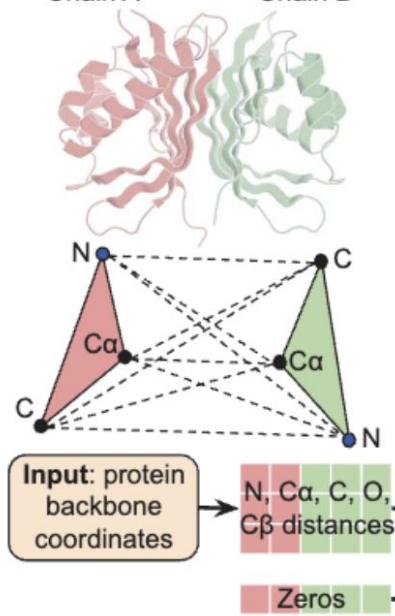
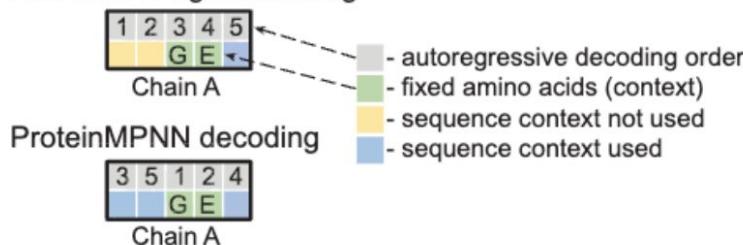
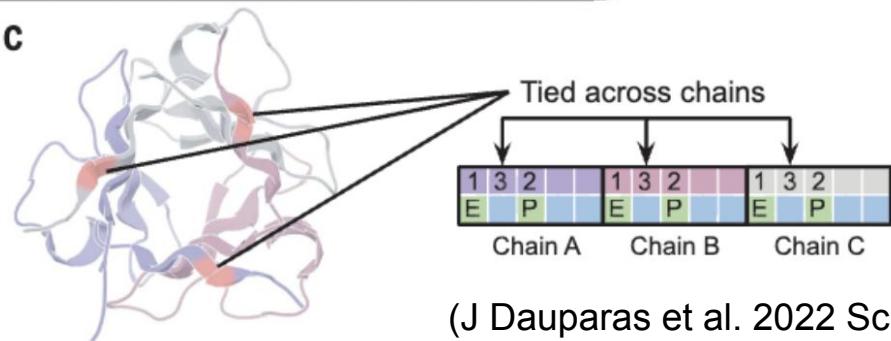
THE ROYAL SWEDISH ACADEMY OF SCIENCES

Live Demo

- [PDB](#)
- [uniprot](#)
- [ColabFold](#)
- [ProteinMPNN](#)
- [RFdiffusion](#)

What about the inverse folding?



A Chain A Chain B**B** Fixed left to right decoding**C**

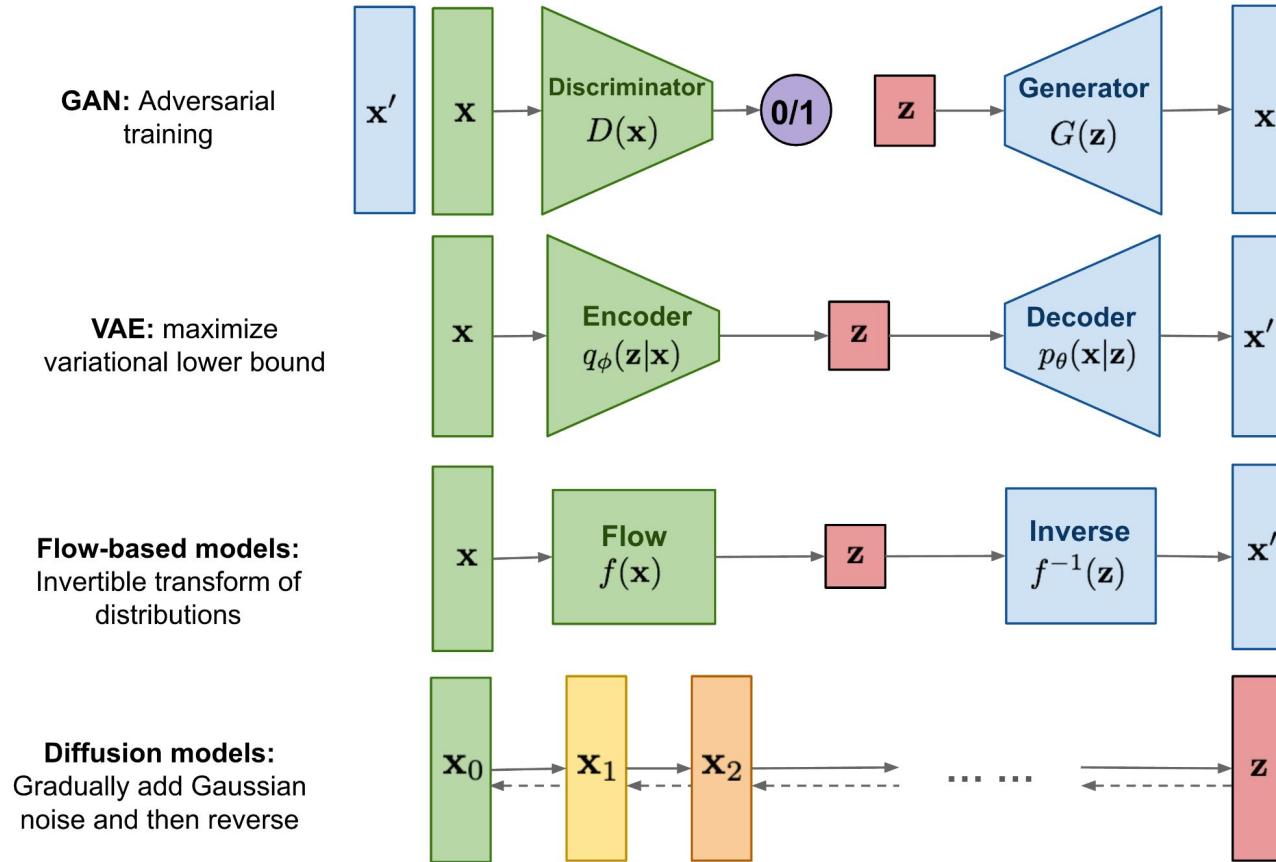
Live Demo

- [PDB](#)
- [uniprot](#)
- [ColabFold](#)
- [ProteinMPNN](#)
- [RFdiffusion](#)

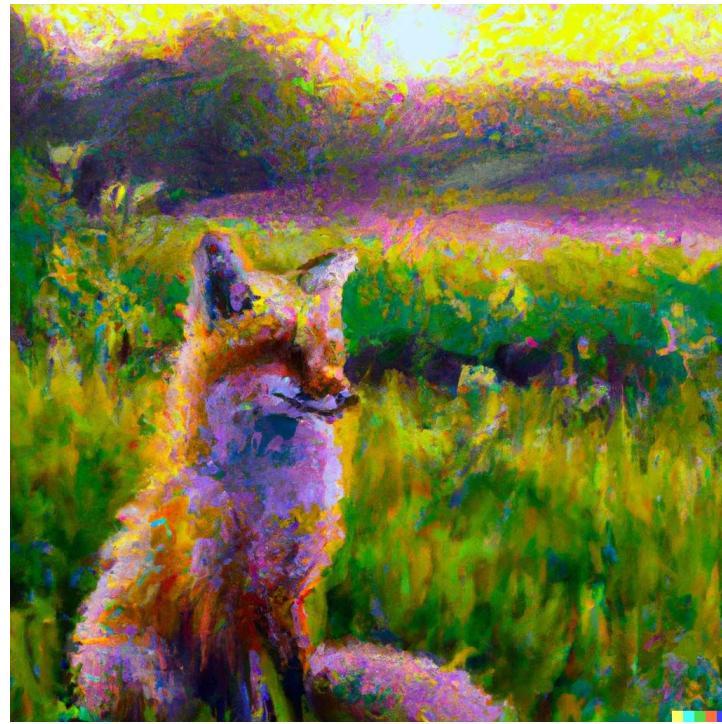
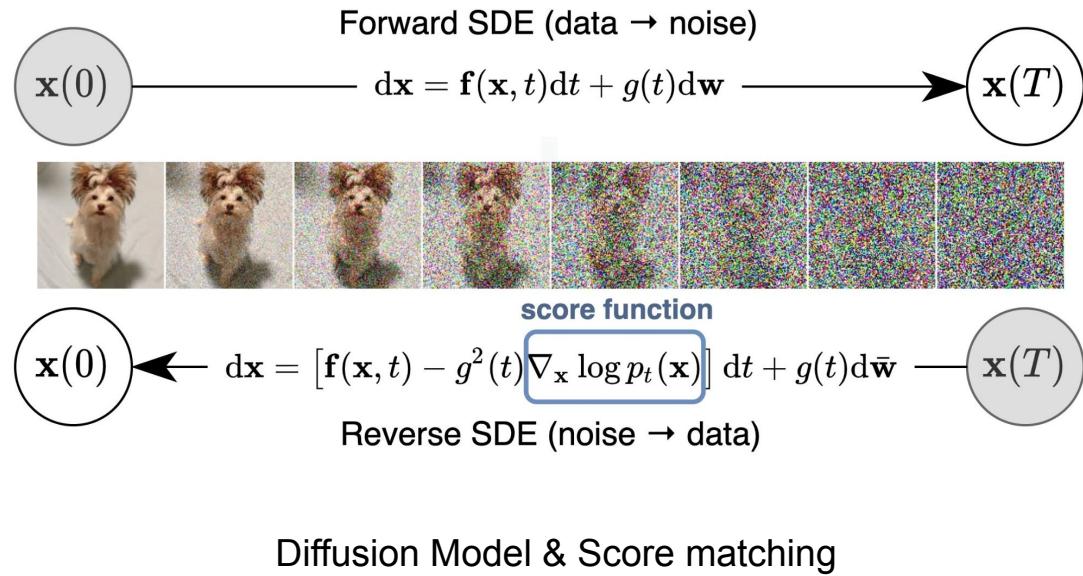
Why do we need ML for protein design

- Even if we could actually predict the structure of protein, the possible amino acid sequence is more than 20^{300} .
- **Generative model** for protein is needed

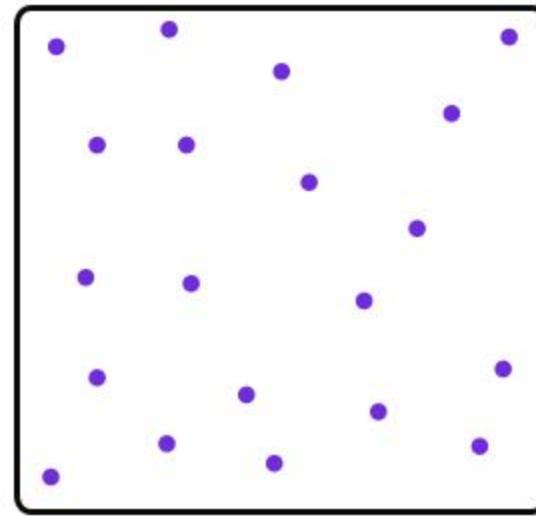
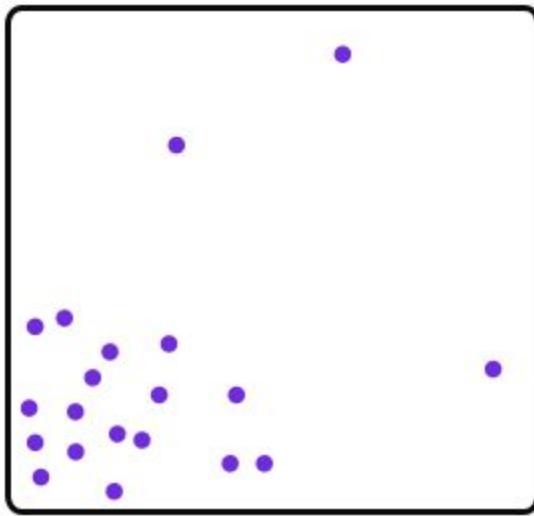
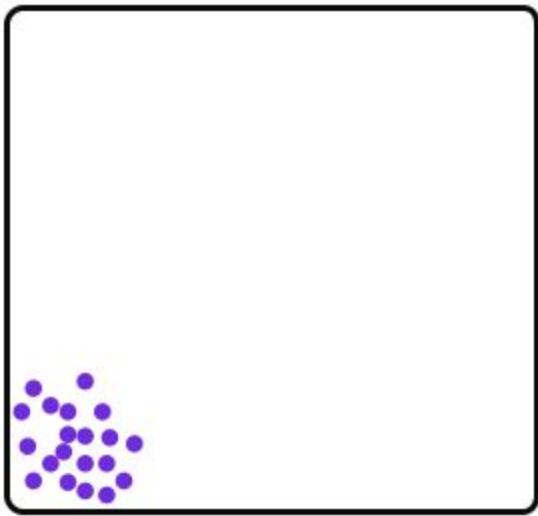
Family of generative models



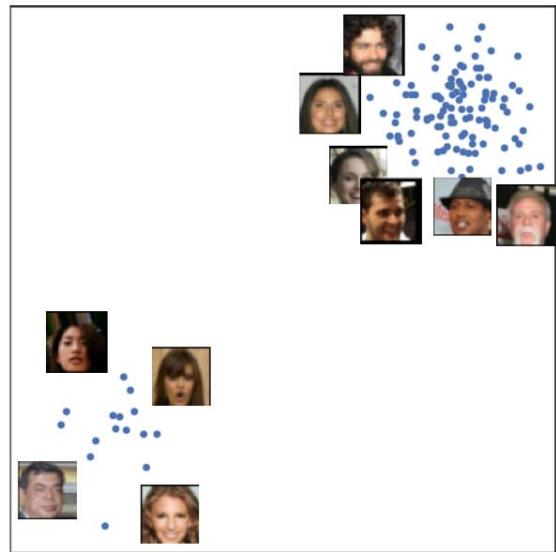
Diffusion Model (& Score matching)



"A painting of a fox sitting in a field at sunrise in the style of Claude Monet"

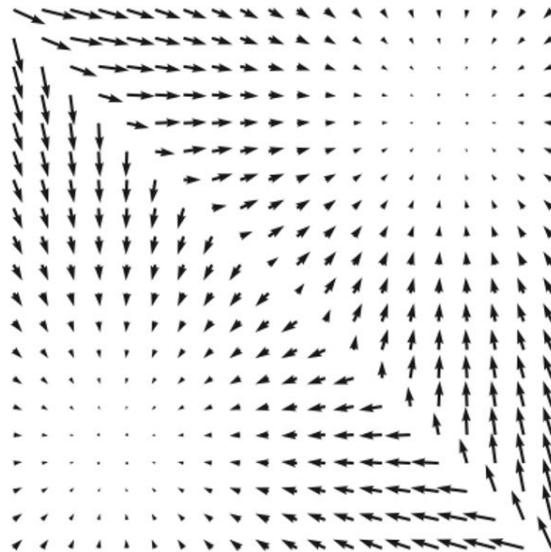


time



$$\{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N\} \stackrel{\text{i.i.d.}}{\sim} p(\mathbf{x})$$

score
matching



$$\mathbf{s}_\theta(\mathbf{x}) \approx \nabla_{\mathbf{x}} \log p(\mathbf{x})$$

Langevin
dynamics



New samples

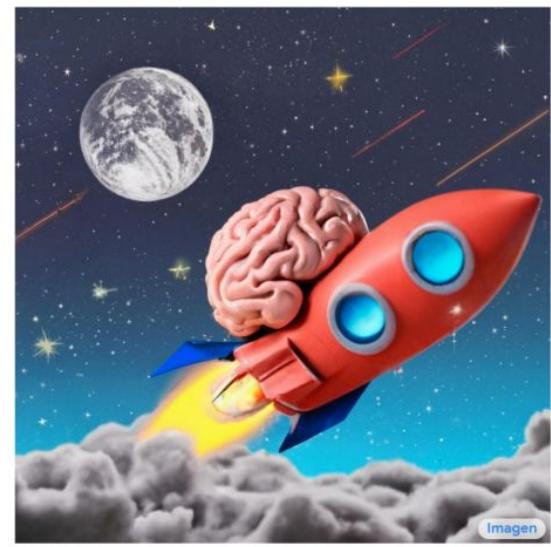
Example from Imagen



“A photo of a Shiba Inu dog with a backpack riding a bike. It is wearing sunglasses and a beach hat.”

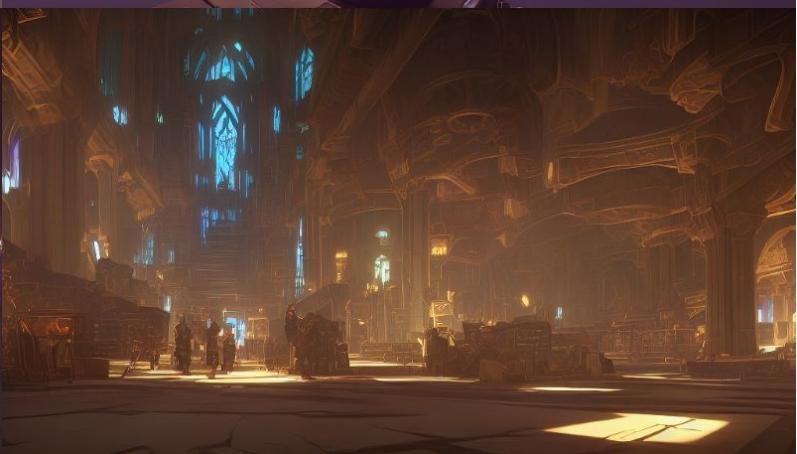
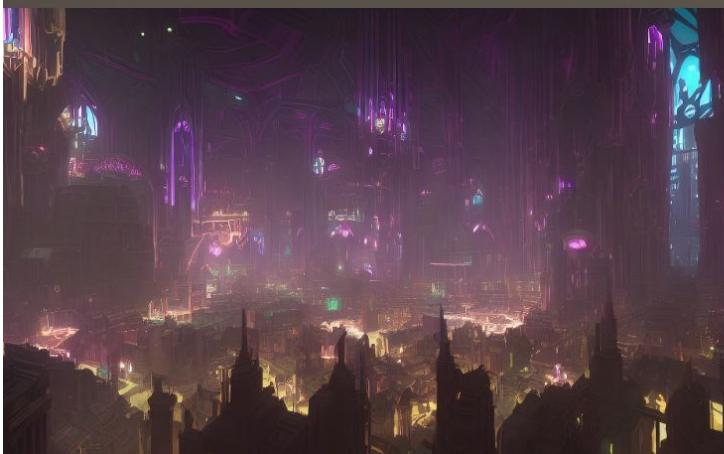


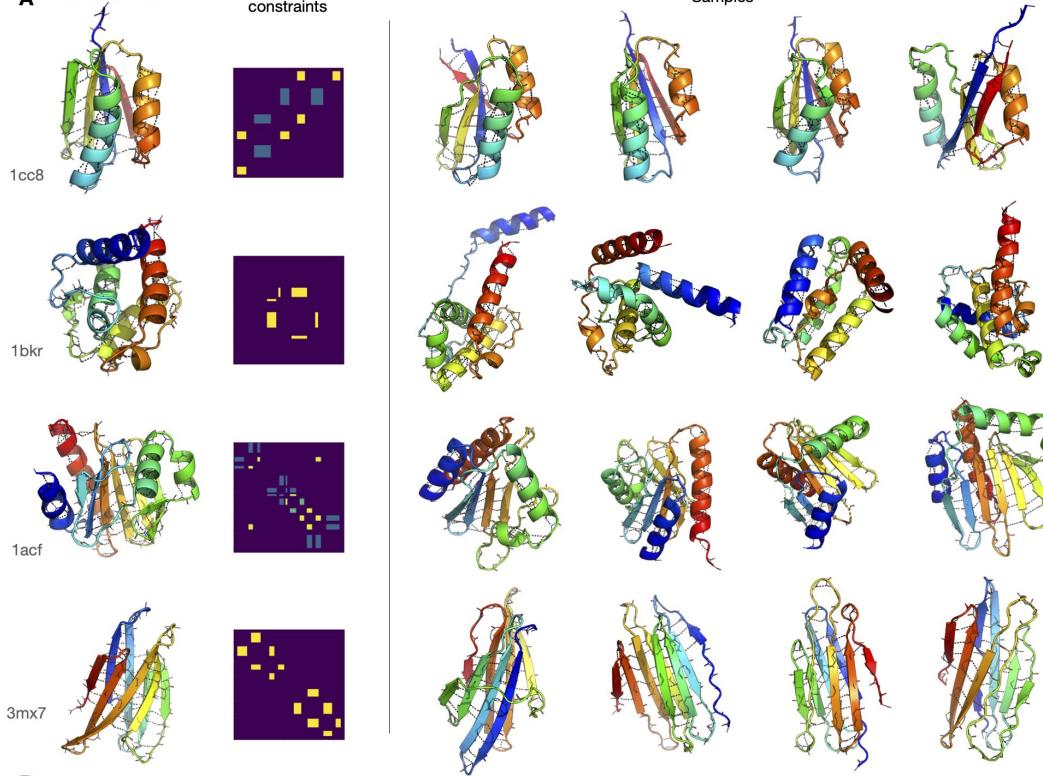
“A blue jay standing on a large basket of rainbow macarons”



“A brain riding a rocketship heading towards the moon”

Example of Stable diffusion

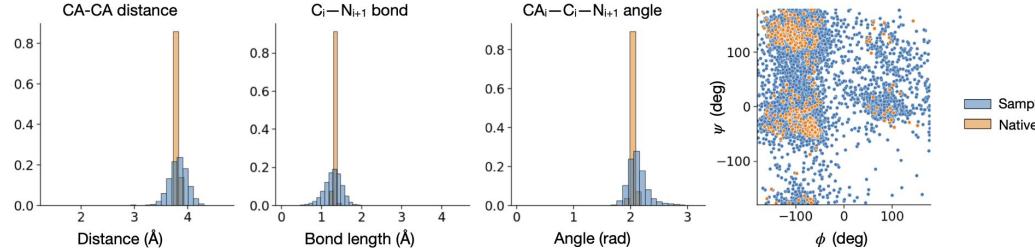


A Native structureBlock adjacency
constraints

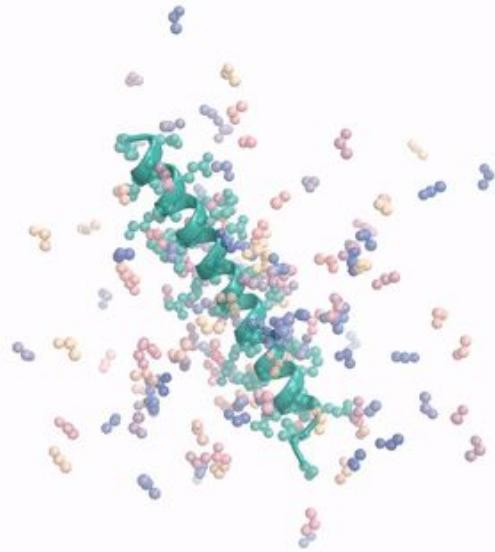
Namrata Anand

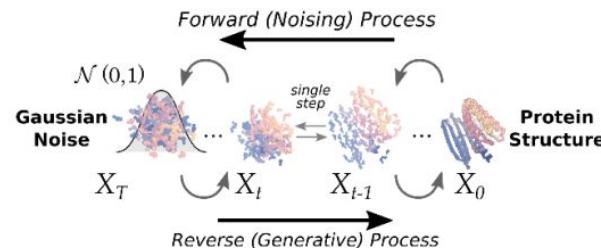
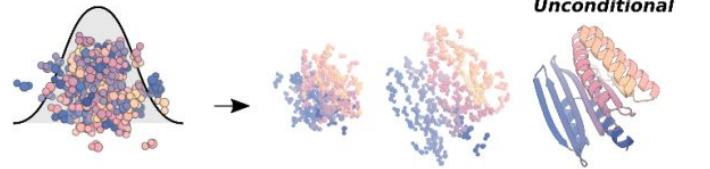
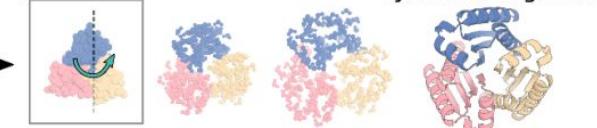
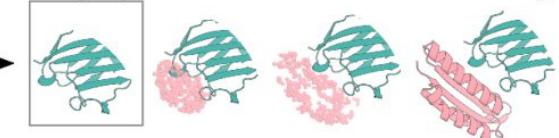
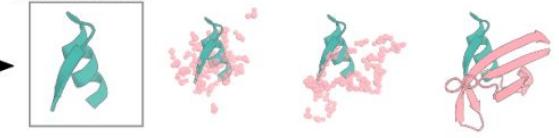
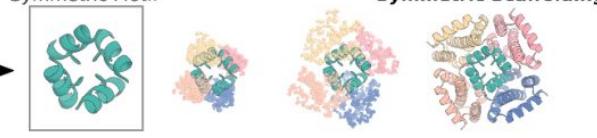
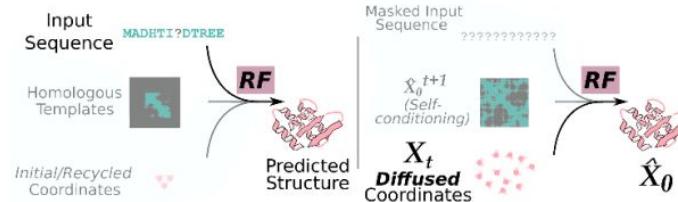
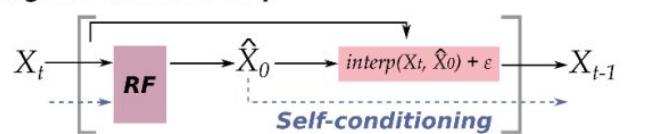


Tudor Achim

B(arXiv: [2205.15019](https://arxiv.org/abs/2205.15019))

RFdiffusion: diffusion model for protein



A**Diffusion Model****B****Unconditional****Symmetric Noise****Binding Target****Functional Motif****Symmetric Motif****RoseTTAFold****Single RFdiffusion step****Figure 1: RFdiffusion is a denoising diffusion probabilistic model with RoseTTAFold**

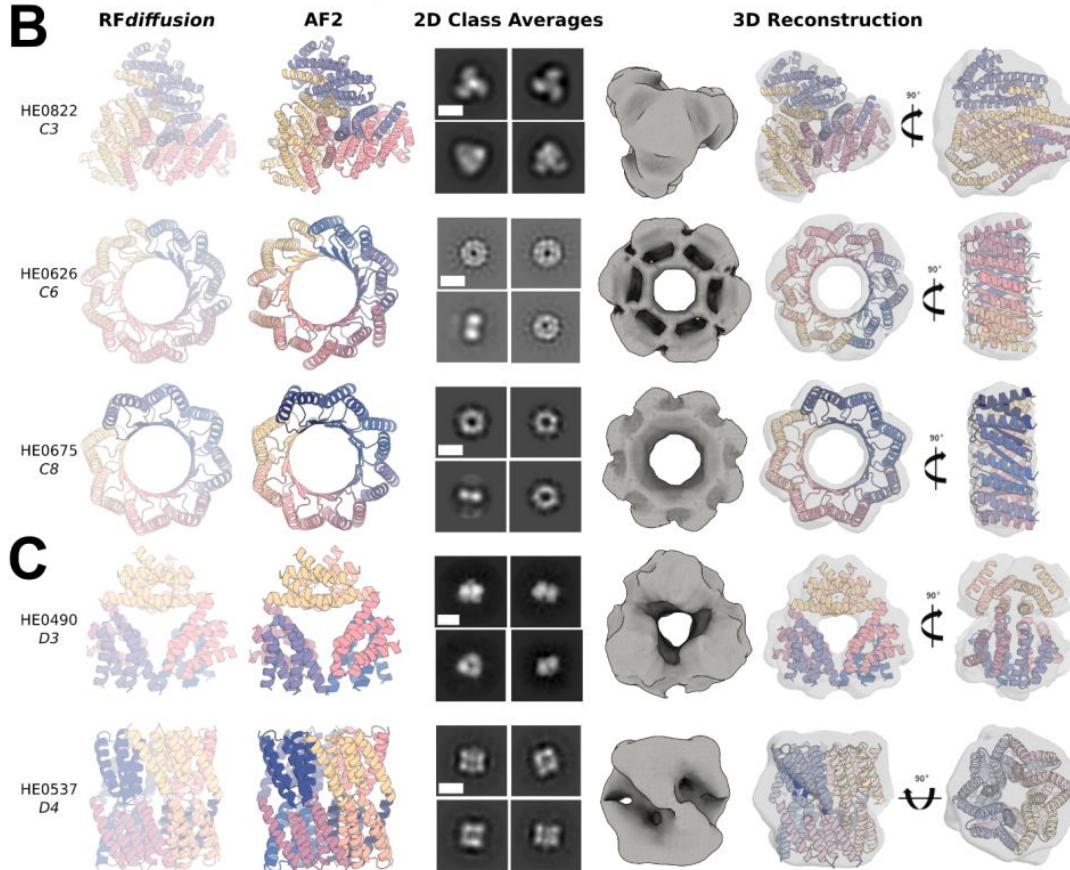
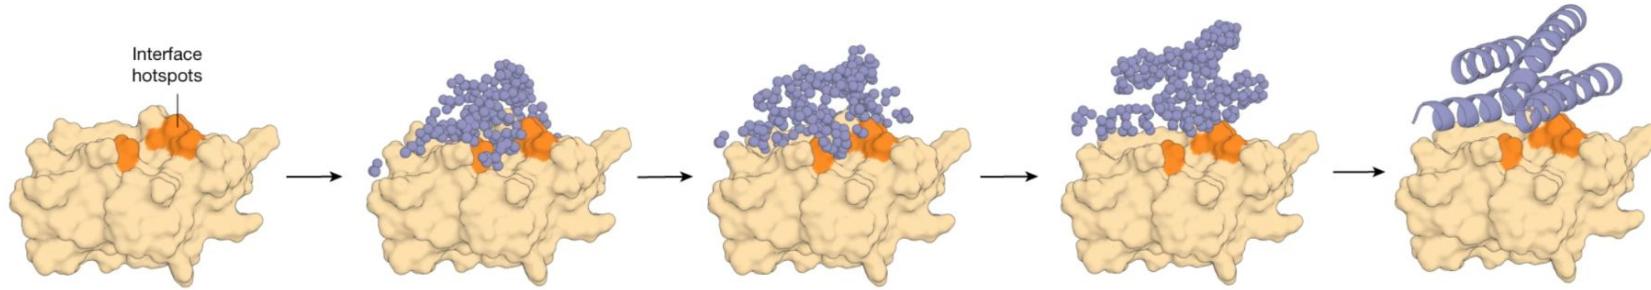
A**B****C**

Fig. 6: De novo design of protein-binding proteins.

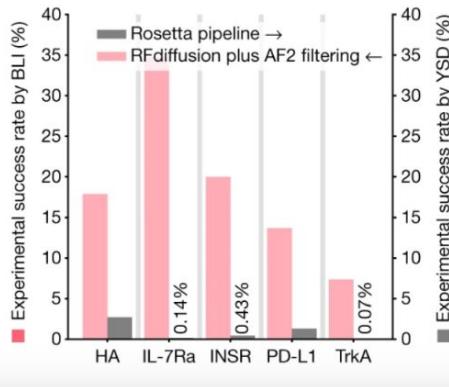
From: [De novo design of protein structure and function with RFdiffusion](#)

a

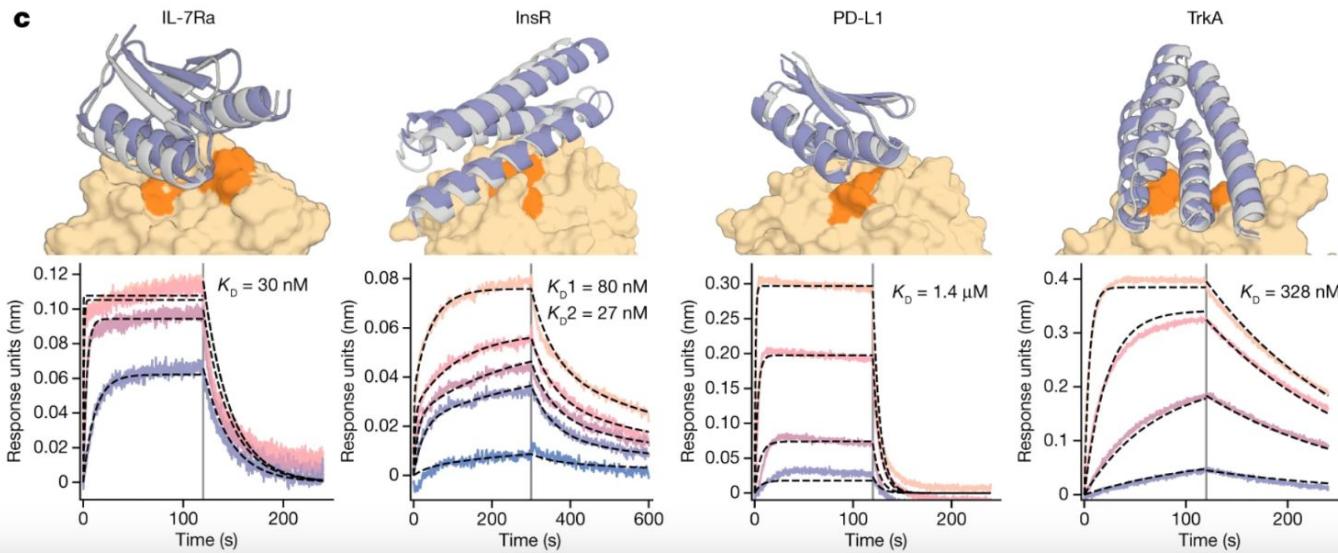


b

RFdiffusion plus AF2 filtering has orders-of-magnitude higher experimental success rates than previous methods



c



De novo designed proteins neutralize lethal snake venom toxins

<https://doi.org/10.1038/s41586-024-08393-x>

Received: 8 May 2024

Accepted: 13 November 2024

Published online: 15 January 2025

Open access

 Check for updates

Susana Vázquez Torres^{1,2,3}, Melisa Benard Valle⁴, Stephen P. Mackessy⁵, Stefanie K. Menzies^{6,7,8}, Nicholas R. Casewell^{6,7}, Shirin Ahmadi⁴, Nick J. Burlet⁴, Edin Muratspahić^{1,2}, Isaac Sappington^{1,2,3}, Max D. Overath⁴, Esperanza Rivera-de-Torre⁴, Jann Ledergerber⁴, Andreas H. Laustsen⁴, Kim Boddum⁹, Asim K. Bera^{1,2}, Alex Kang^{1,2}, Evans Brackenbrough^{1,2}, Iara A. Cardoso⁶, Edouard P. Crittenden⁶, Rebecca J. Edge¹⁰, Justin Decarreau^{1,2}, Robert J. Ragotte^{1,2}, Arvind S. Pillai^{1,2}, Mohamad Abedi^{1,2}, Hannah L. Han^{1,2}, Stacey R. Gerben^{1,2}, Analisa Murray^{1,2}, Rebecca Skotheim^{1,2}, Lynda Stuart^{1,2}, Lance Stewart^{1,2}, Thomas J. A. Fryer^{4,11}, Timothy P. Jenkins⁴ & David Baker^{1,2,12}



Snakebite envenoming remains a devastating and neglected tropical disease, claiming over 100,000 lives annually and causing severe complications and long-lasting disabilities for many more^{1,2}. Three-finger toxins (3FTx) are highly toxic components of elapid snake venoms that can cause diverse pathologies, including severe tissue damage³ and inhibition of nicotinic acetylcholine receptors, resulting in life-threatening neurotoxicity⁴. At present, the only available treatments for snakebites consist of polyclonal antibodies derived from the plasma of immunized animals, which have high cost and limited efficacy against 3FTx^{5–7}. Here we used deep learning methods to de novo design proteins to bind short-chain and long-chain α -neurotoxins and cytotoxins from the 3FTx family. With limited experimental screening, we obtained protein designs with remarkable thermal stability, high binding affinity and near-atomic-level agreement with the computational models. The designed proteins effectively neutralized all three 3FTx subfamilies in vitro and protected mice from a lethal neurotoxin challenge. Such potent, stable and readily manufacturable toxin-neutralizing proteins could provide the basis for safer, cost-effective and widely accessible next-generation antivenom therapeutics. Beyond snakebite, our results highlight how computational design could help democratize therapeutic discovery, particularly in resource-limited settings, by substantially reducing costs and resource requirements for the development of therapies for neglected tropical diseases.



Live Demo

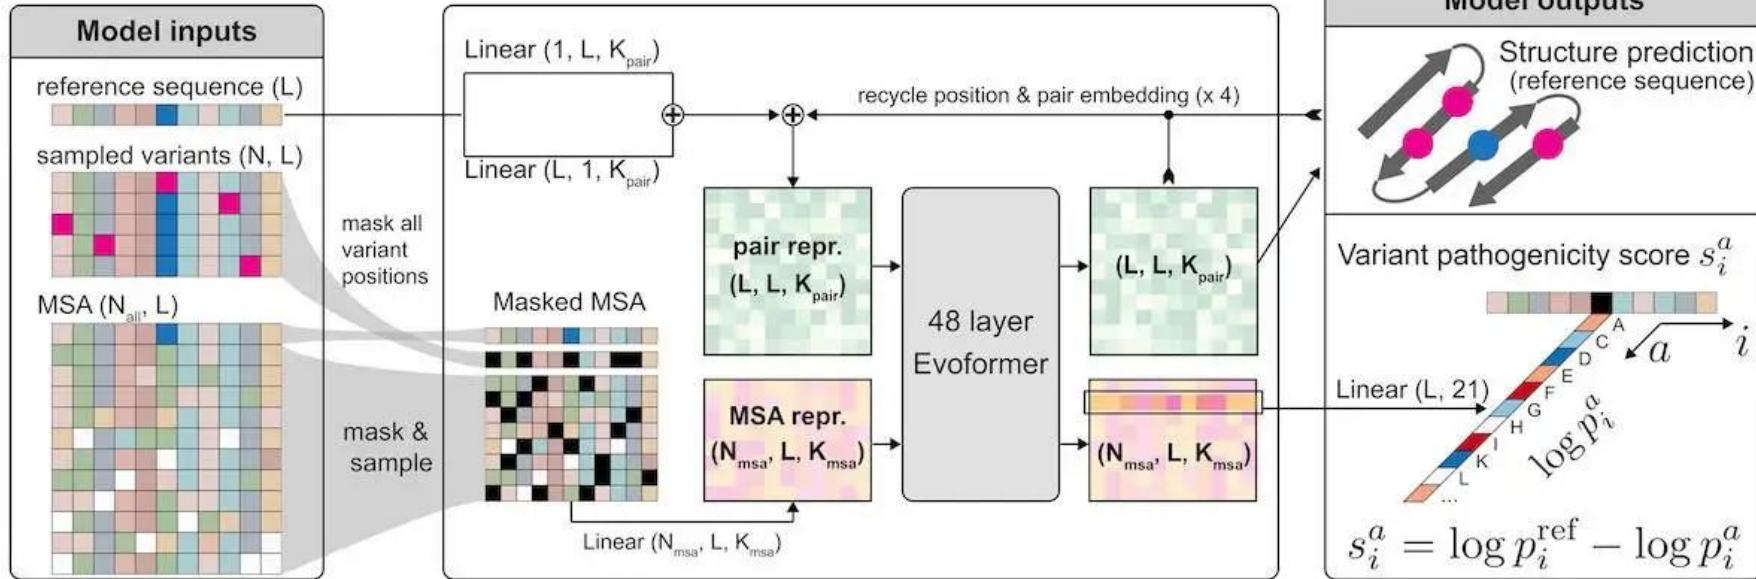
- [PDB](#)
- [uniprot](#)
- [ColabFold](#)
- [ProteinMPNN](#)
- [RFdiffusion](#)



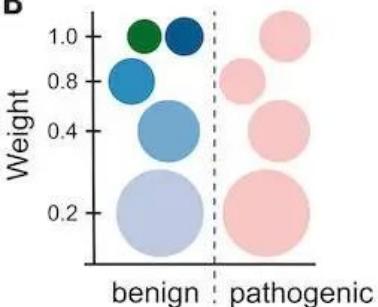
AlphaMissense

AlphaFold Based Tool to
Predict Mutations Causing
Genetic Diseases

A

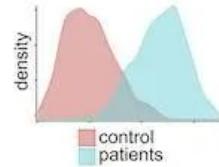
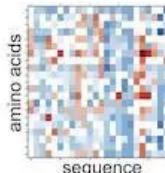


B



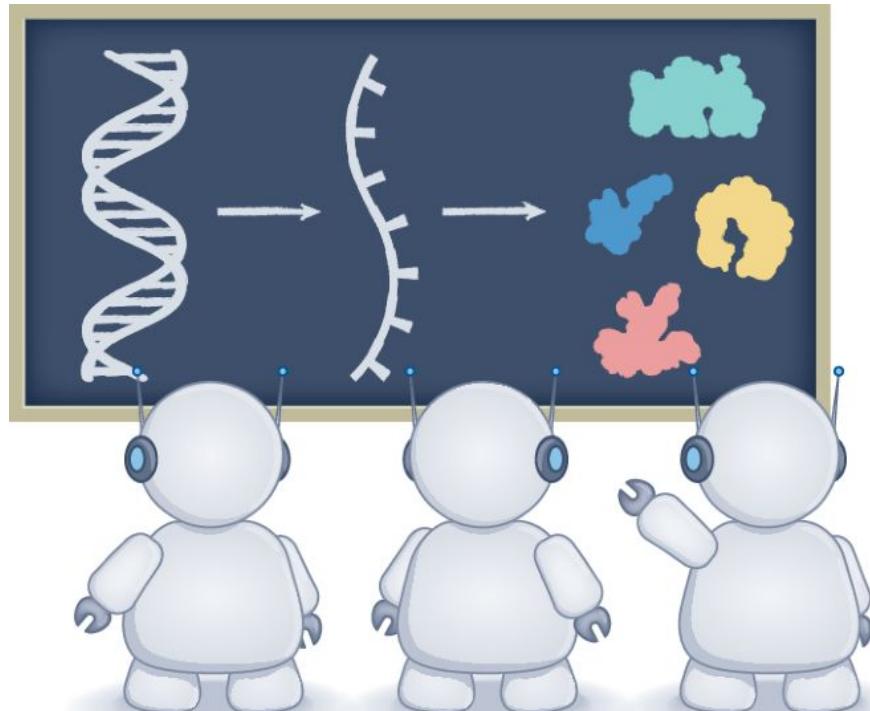
(Jun Cheng et al., Science 2023)

C

Known benign
and pathogenic
variants (ClinVar)De novo
variants of
rare diseaseMultiplexed assay
of variant effect
(ProteinGym)

Where are we at AI for Biology?

- Where are we at for drug discovery?
- Where are we at for protein dynamics?
- Where are we at for genetics?

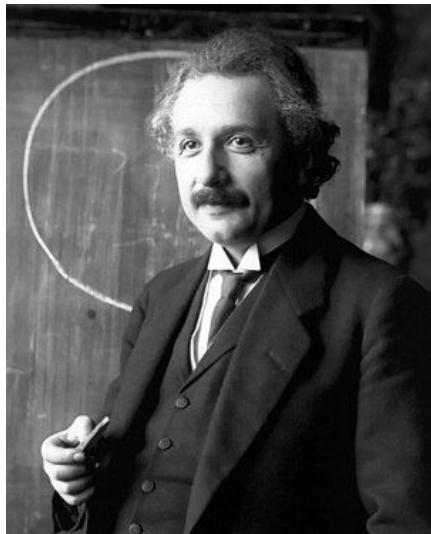


Some useful Resources

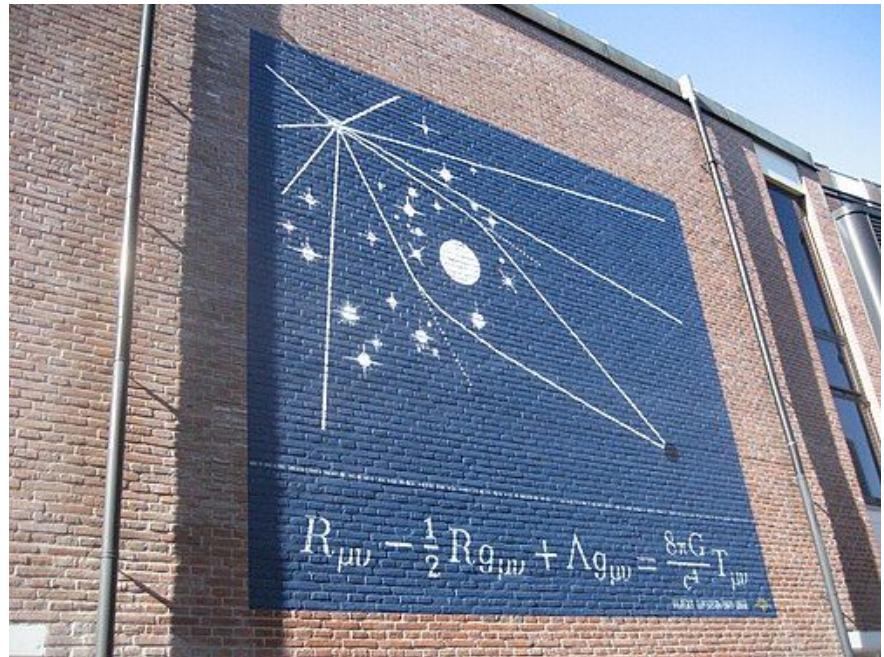
- [PDB](#)
- [uniprot](#)
- [ColabFold](#)
- [ProteinMPNN](#)
- [RFdiffusion](#)
- [Alphafold server](#)
- [Foldseek](#)
- [Kaggle](#)
- [AlphafoldDB](#)

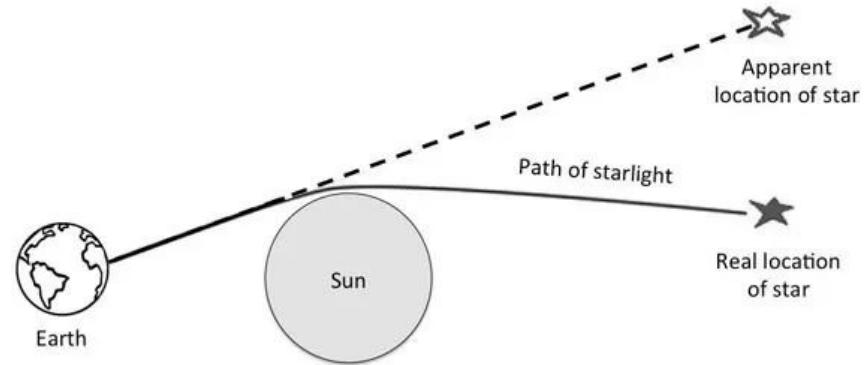


Image Credit: Event Horizon Telescope Collaboration



Albert Einstein





Eddington experiment

Arthur Eddington



Karl Schwarzschild

SCHWARZSCHILD METRIC

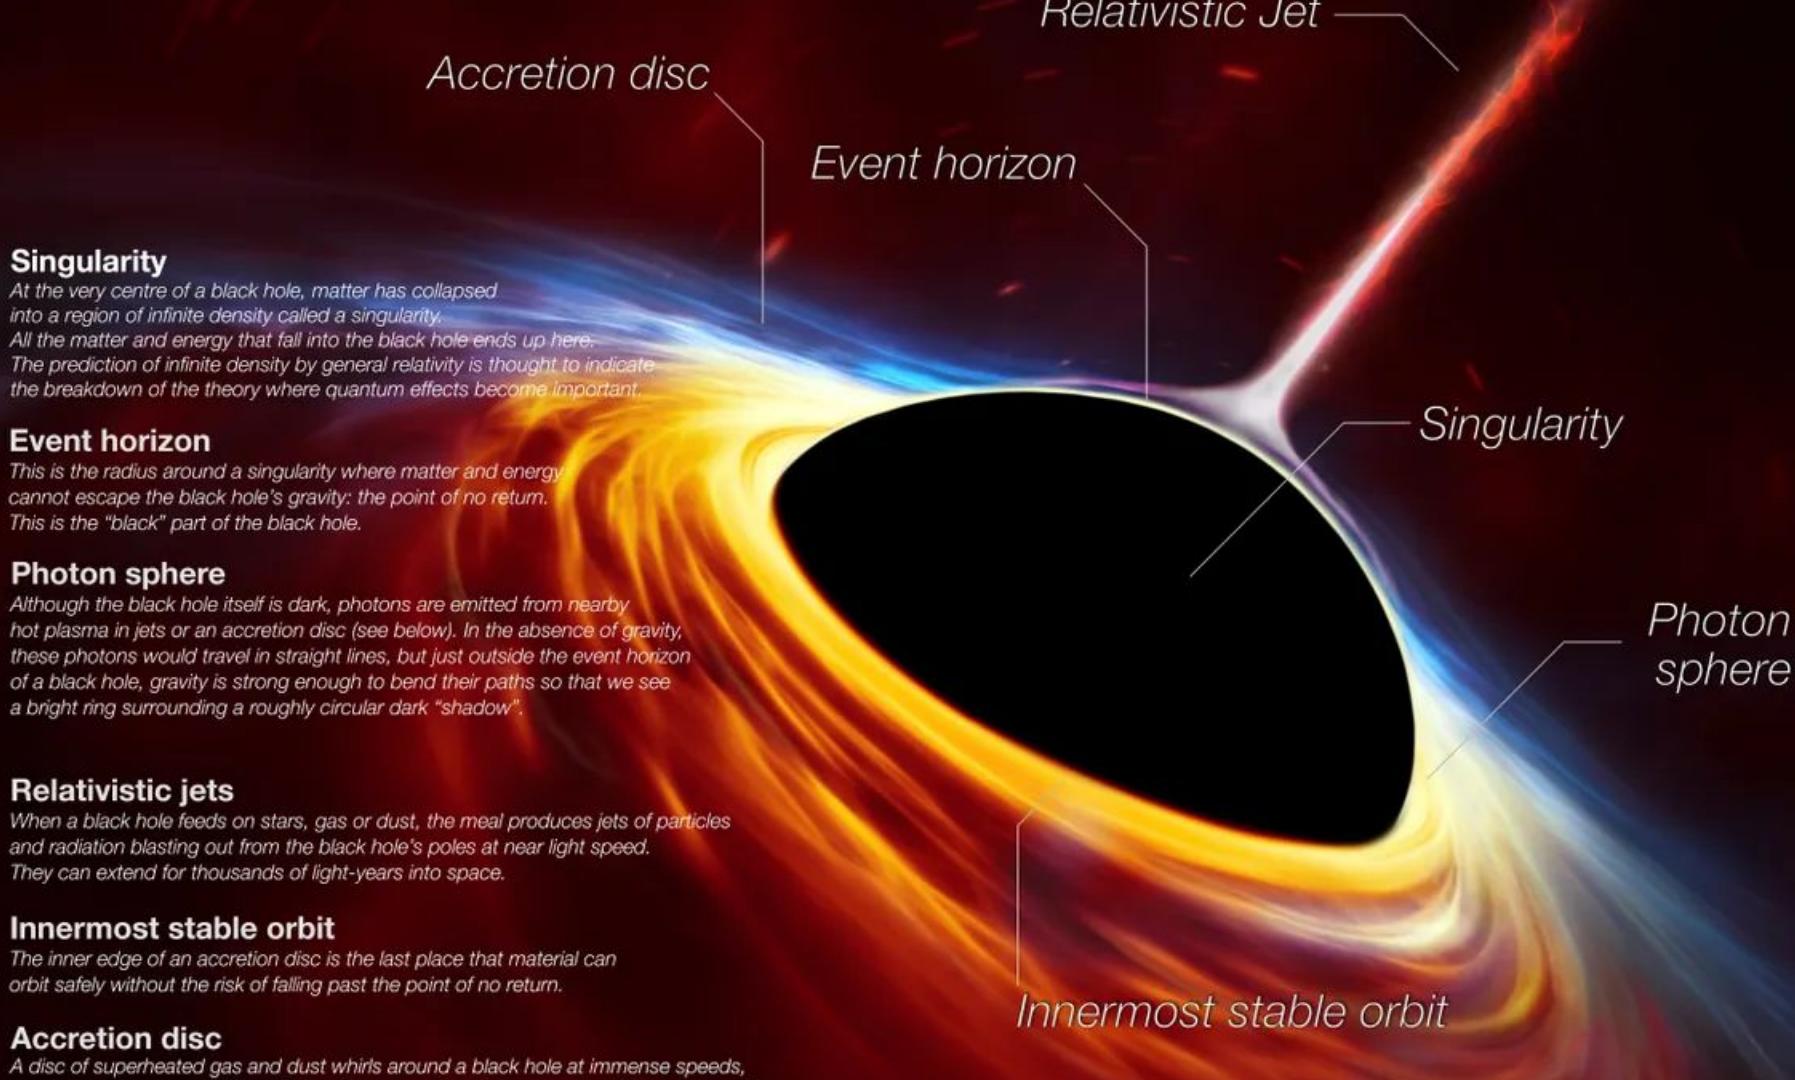
describes the shape of spacetime around a spherical source, where $g_{\mu\nu}$ is the spacetime metric:

$$g_{\mu\nu} = \begin{pmatrix} g_{tt} & g_{tr} & g_{t\theta} & g_{t\varphi} \\ g_{rt} & g_{rr} & g_{r\theta} & g_{r\varphi} \\ g_{\theta t} & g_{\theta r} & g_{\theta\theta} & g_{\theta\varphi} \\ g_{\varphi t} & g_{\varphi r} & g_{\varphi\theta} & g_{\varphi\varphi} \end{pmatrix}$$
$$= \begin{pmatrix} -\left(1 - \frac{2G_NM}{c^2r}\right) & 0 & 0 & 0 \\ 0 & \left(1 - \frac{2G_NM}{c^2r}\right)^{-1} & 0 & 0 \\ 0 & 0 & r^2 & 0 \\ 0 & 0 & 0 & r^2 \sin^2 \theta \end{pmatrix}$$

coordinates: $(x^0, x^1, x^2, x^3) = (ct, r, \theta, \varphi)$



Karl Schwarzschild



SIZE COMPARISON:
THE M87 BLACK HOLE
AND
OUR SOLAR SYSTEM

EHT BLACK HOLE IMAGE
SOURCE: NSF



~ 40 uas

HST best resolution ~ 0.04 as
ALMA best resolution ~ 0.01 as

Black Hole Observation

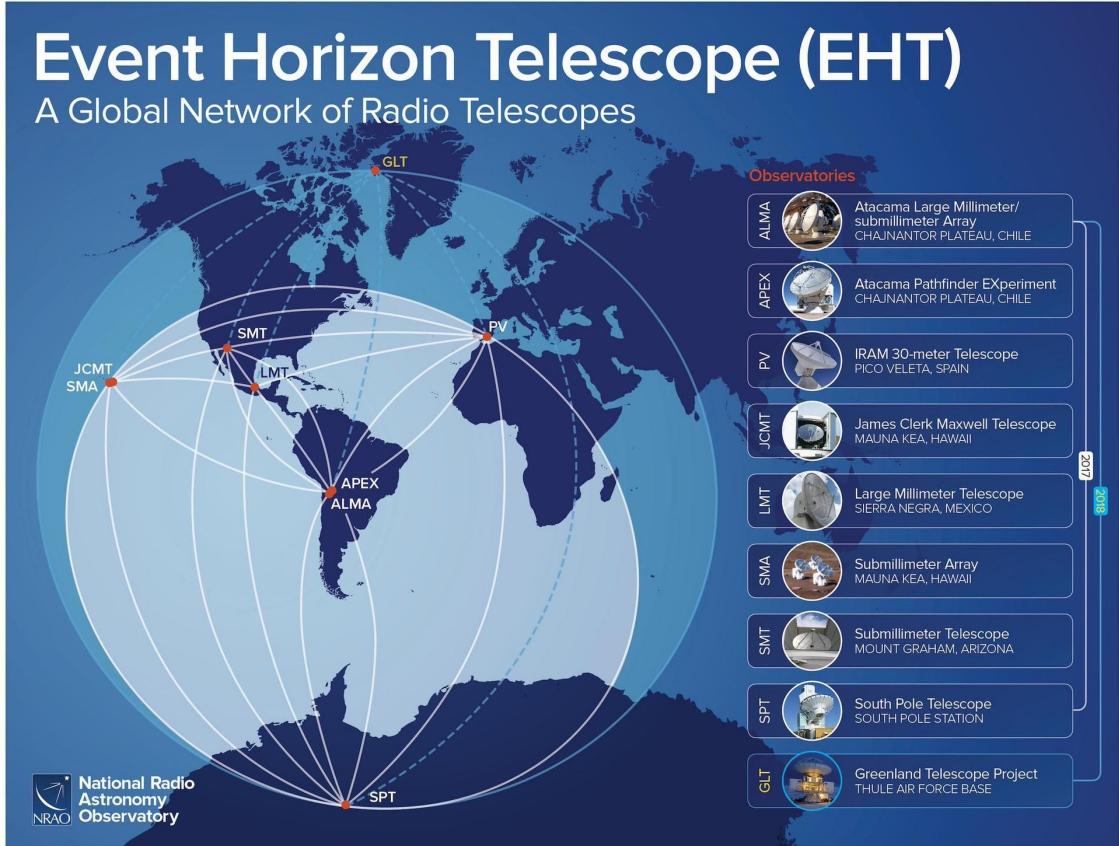
- **Wavelength (230GHz, ~1.3mm)** ✓
- Right size of the telescope ~ (13 Million meters = diameter of the earth)



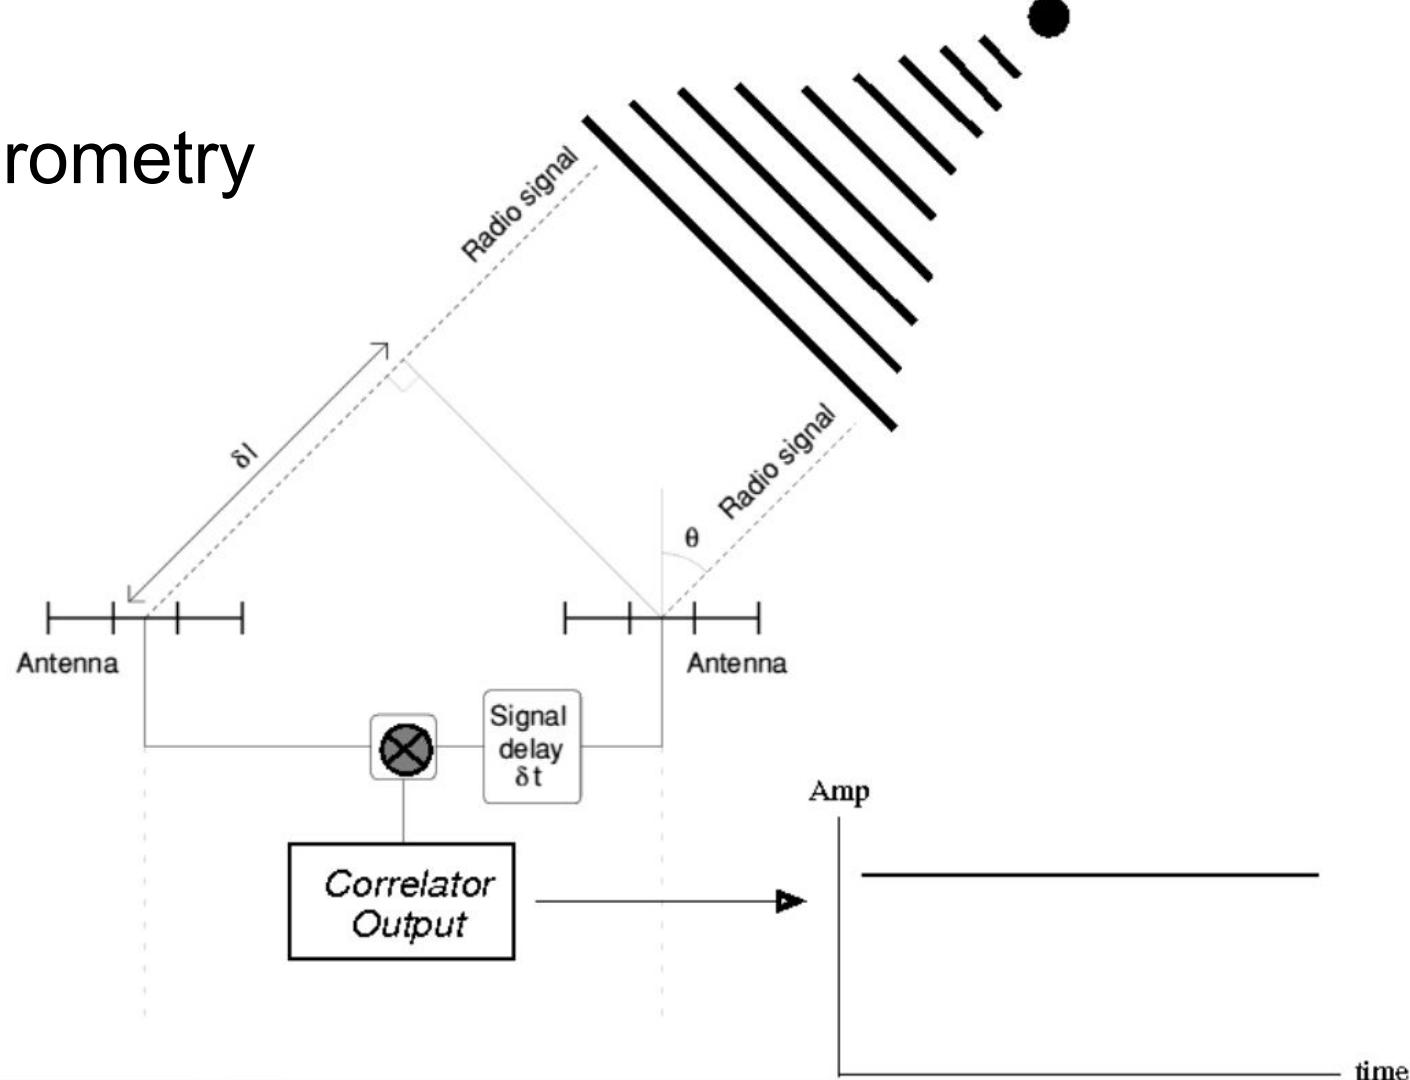
ALMA Telescope Credits: ESO

Black Hole Observation

- Wavelength (230GHz, ~1.3mm) ✓
- Right size of the telescope ~ (13 Million meters = diameter of the earth) ✓



Radio Interferometry



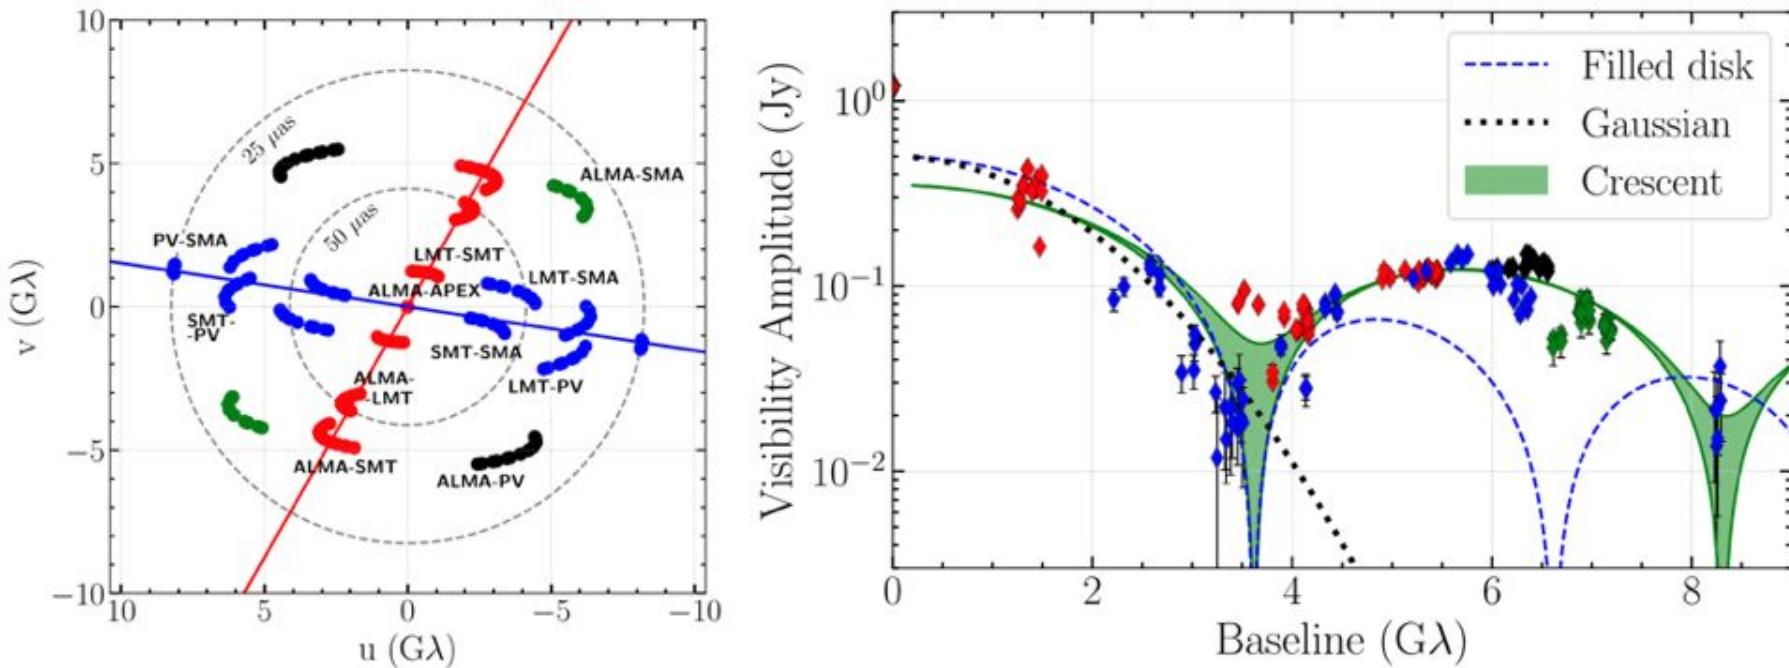


Figure 1. (u , v)-coverage (left panel) and visibility amplitudes (right panel) of M87 for the high-band April 11 data. The (u , v)-coverage has two primary orientations, east–west in blue and north–south in red, with two diagonal fillers at large baselines in green and black. Note that the Large Millimeter Telescope (LMT) and the Submillimeter Telescope (SMT) participate in both orientations, and that the LMT amplitudes are subject to significant gain errors. There is evidence for substantial depressions in the visibility amplitudes at ~ 3.4 G λ and ~ 8.3 G λ . The various lines in the right panel show the expected behavior of (dotted line) a Gaussian, (dashed line) a filled disk, and (green area) a crescent shape along different orientations. The image of M87 does not appear to be consistent with a filled disk or a Gaussian.



Machine Learning

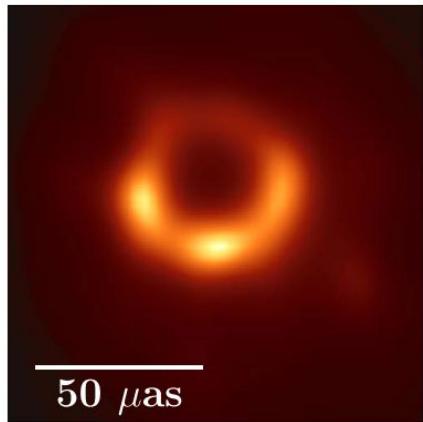


Traditional Method

THE ASTROPHYSICAL JOURNAL LETTERS, 875:L4 (52pp), 2019 April 10

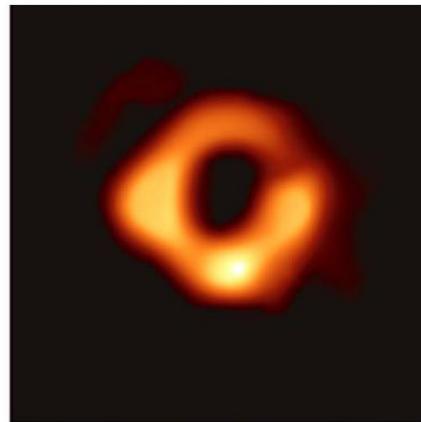
The EHT Collaboration et al.

Team 1 (RML)



50 μ as

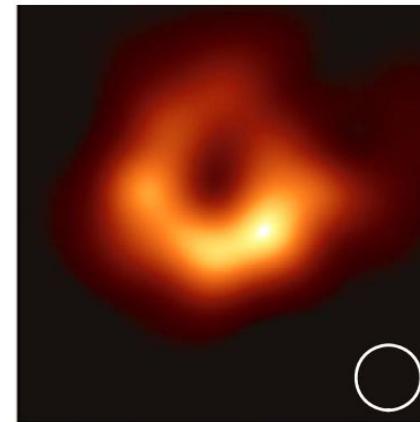
Team 2 (RML)



0.0 2.5 5.0

Brightness Temperature (10^9 K)

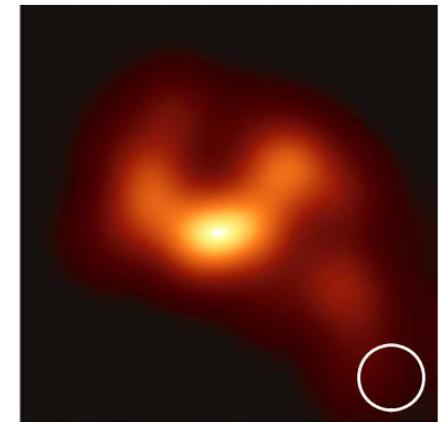
Team 3 (CLEAN)



0 2 4



Team 4 (CLEAN)



0 2 4



Figure 4. The first EHT images of M87, blindly reconstructed by four independent imaging teams using an early, engineering release of data from the April 11 observations. These images all used a single polarization (LCP) rather than Stokes I , which is used in the remainder of this Letter. Images from Teams 1 and 2 used RML methods (no restoring beam); images from Teams 3 and 4 used CLEAN (restored with a circular 20 μ as beam, shown in the lower right). The images all show similar morphology, although the reconstructions show significant differences in brightness temperature because of different assumptions regarding the total compact flux density (see Table 2) and because restoring beams are applied only to CLEAN images.

Image reconstruction (for EHT)

Regularized Maximum Likelihood (RML)

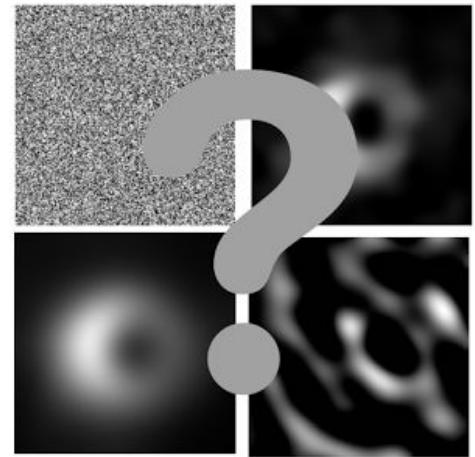
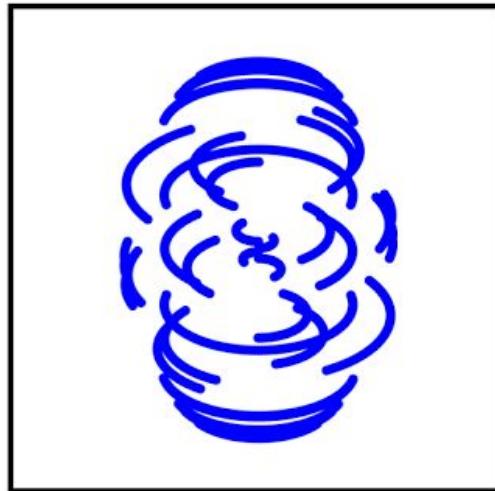
- Used for reconstructing M87* image for EHT
- Prior and forward model (image -> visibility) needed
- Combining with **Normalizing Flow** -> **Deep Probabilistic Imaging**

Clean (Deconvolution)

- Used for reconstructing M87* image for EHT
- Also used for most of the Images Reconstruction for Radio Interferometry
- No prior “hard coded”, but require lots of human expertise and experience on hyperparameter tuning...
- Hard to evaluate uncertainty

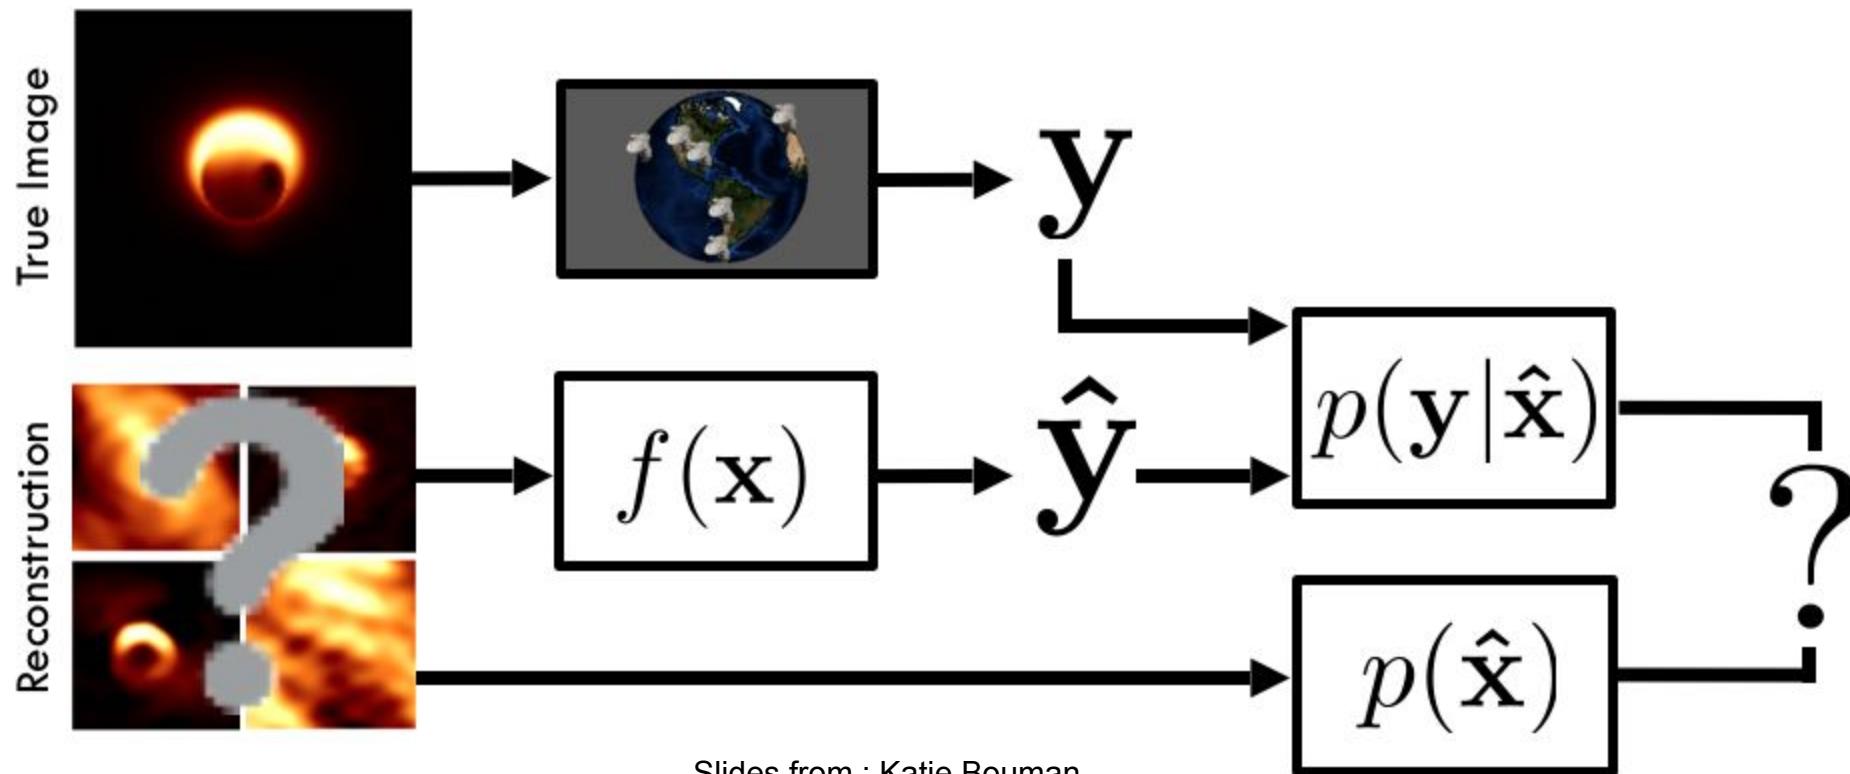


Measurements

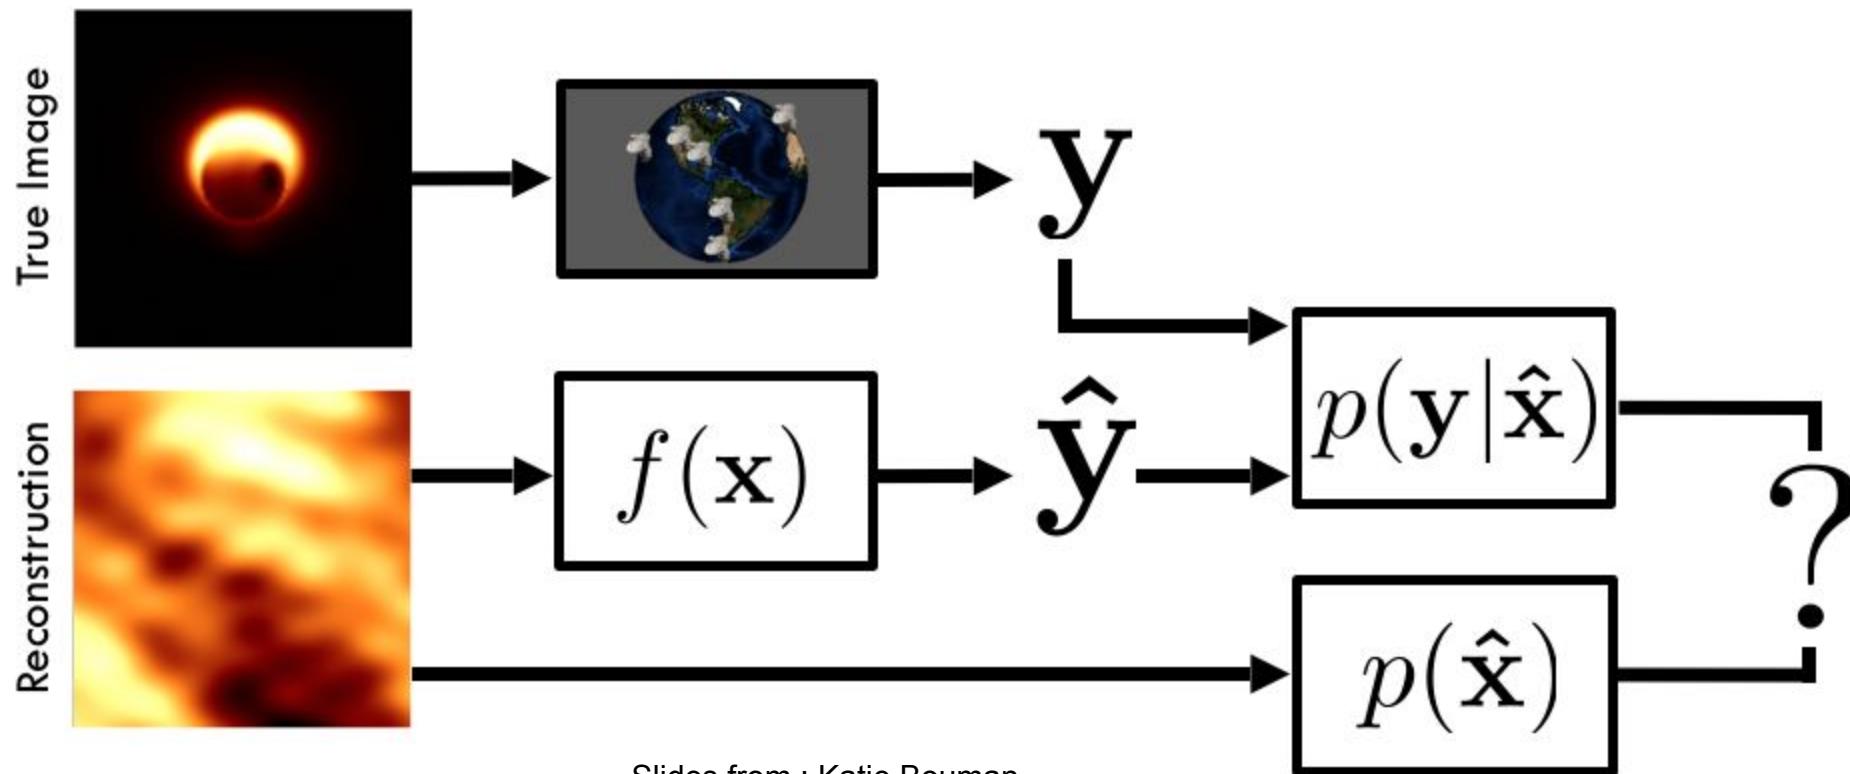


Infinite Number
of Possibilities

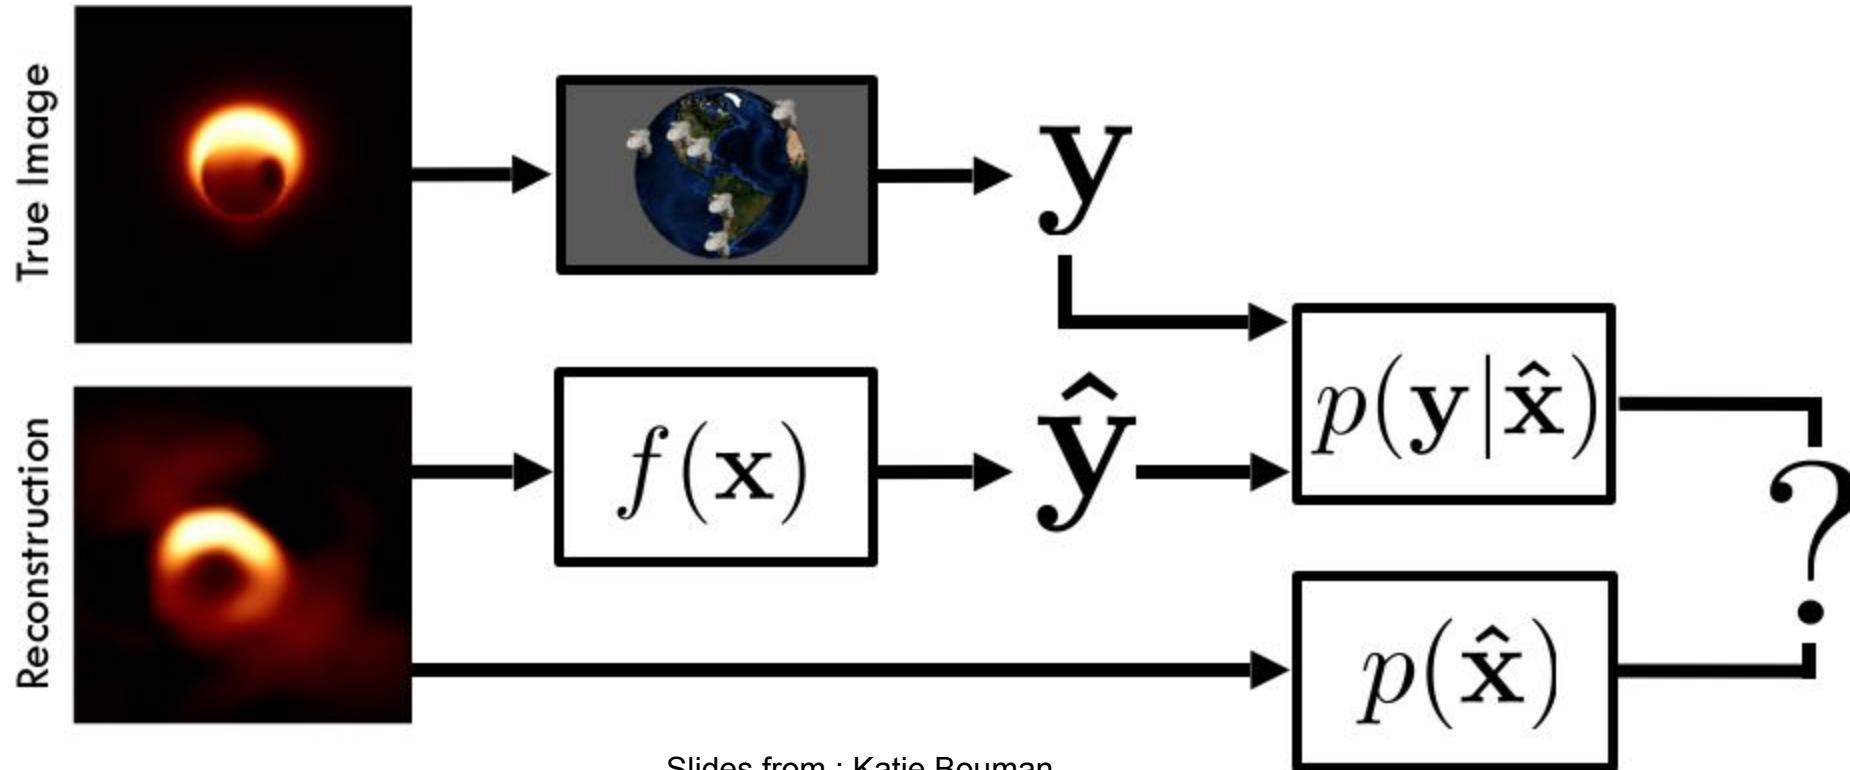
Bayesian Model Inversion



Bayesian Model Inversion



Bayesian Model Inversion



Best Image



$$\hat{\mathbf{x}}_{\text{MAP}} = \operatorname{argmax}_{\mathbf{x}} [\log p(\mathbf{y}|\mathbf{x}) + \log p(\mathbf{x})]$$

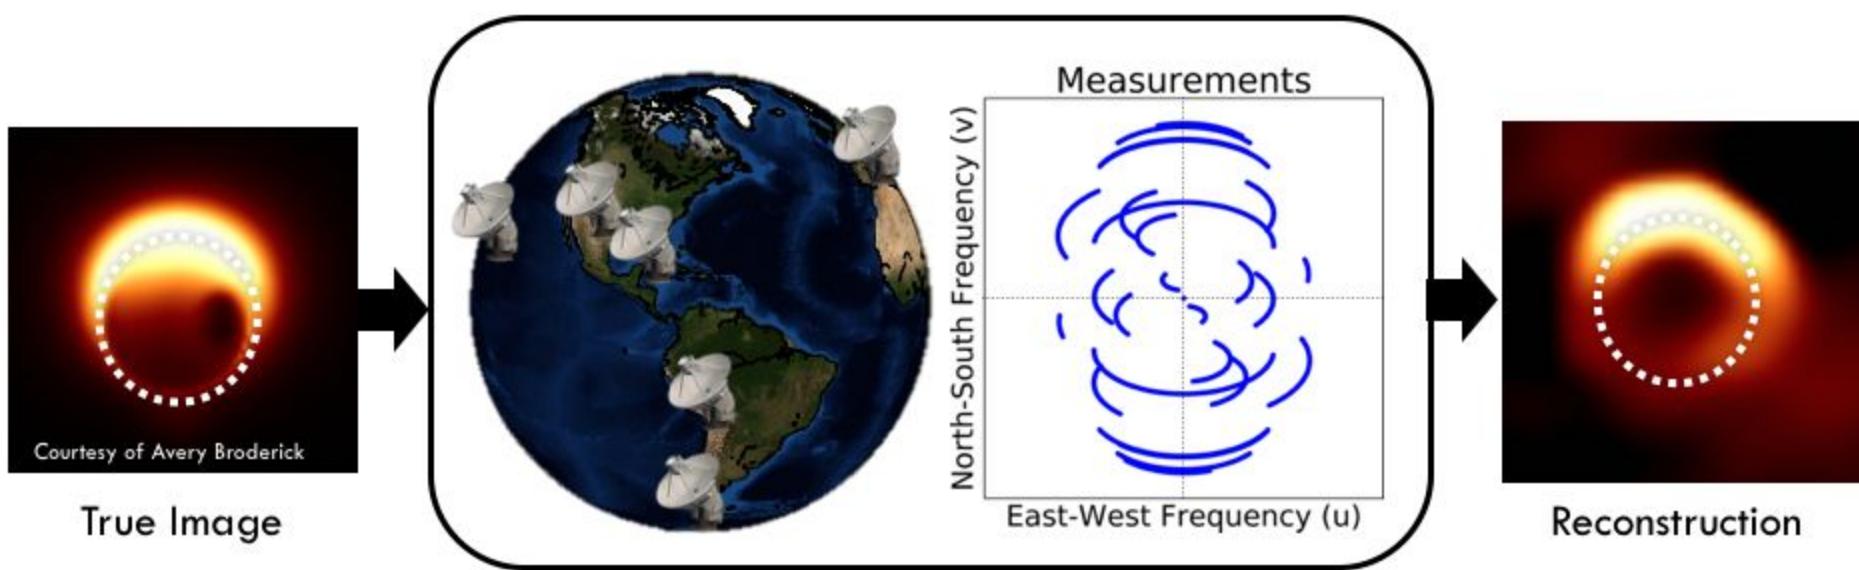
Likelihood



Prior



Simulated EHT Image Reconstruction



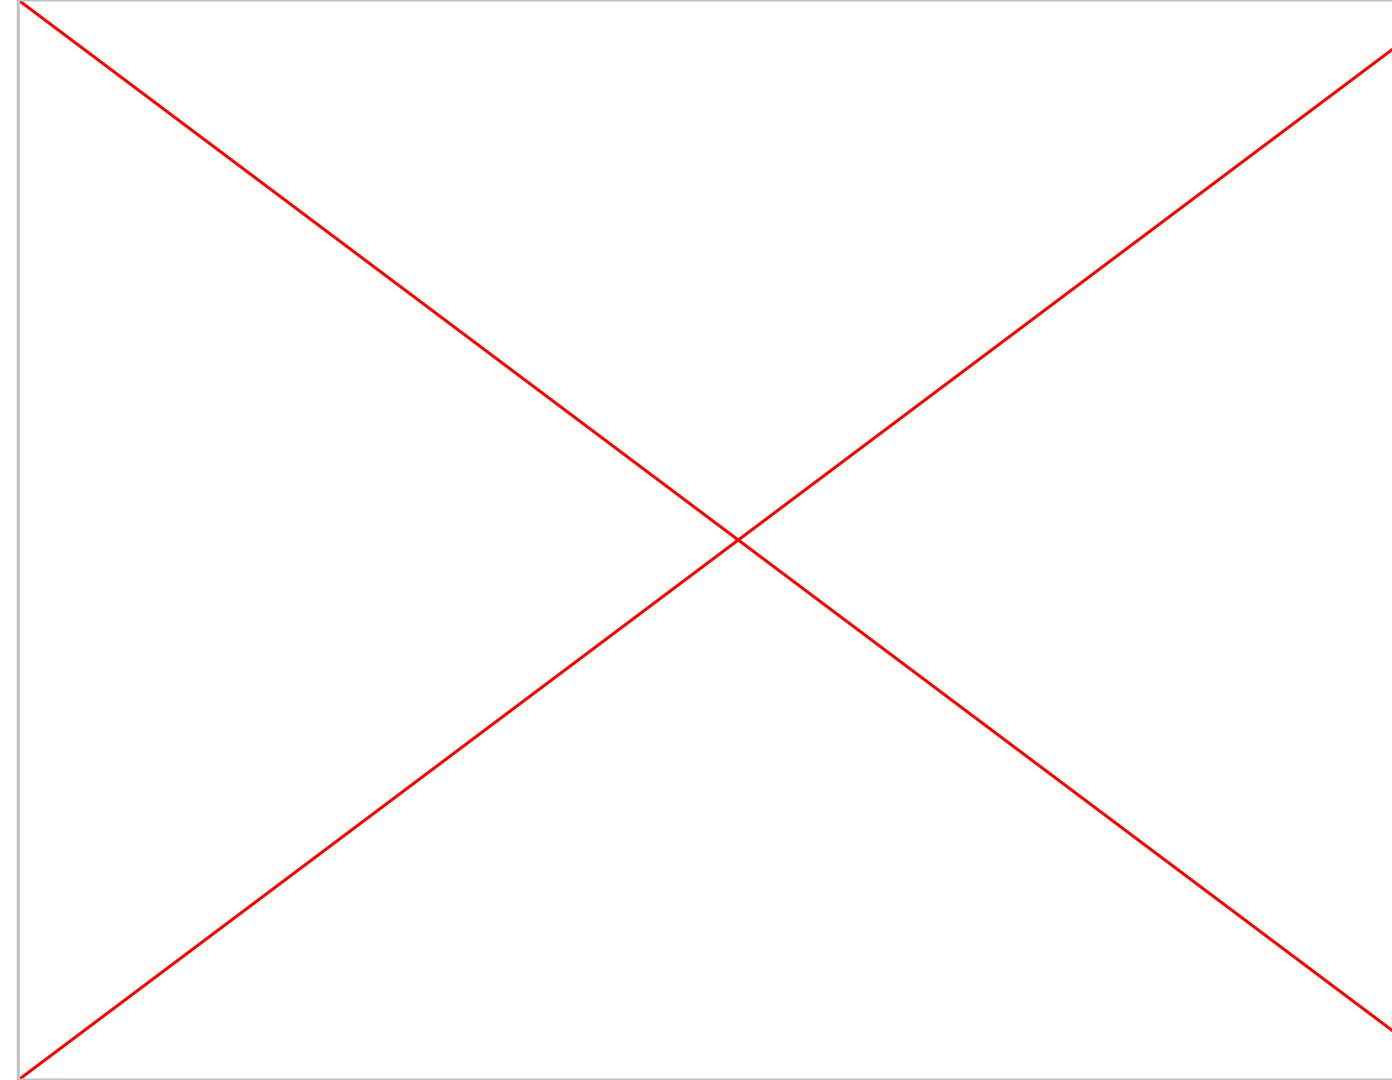
EHT-imaging Demo



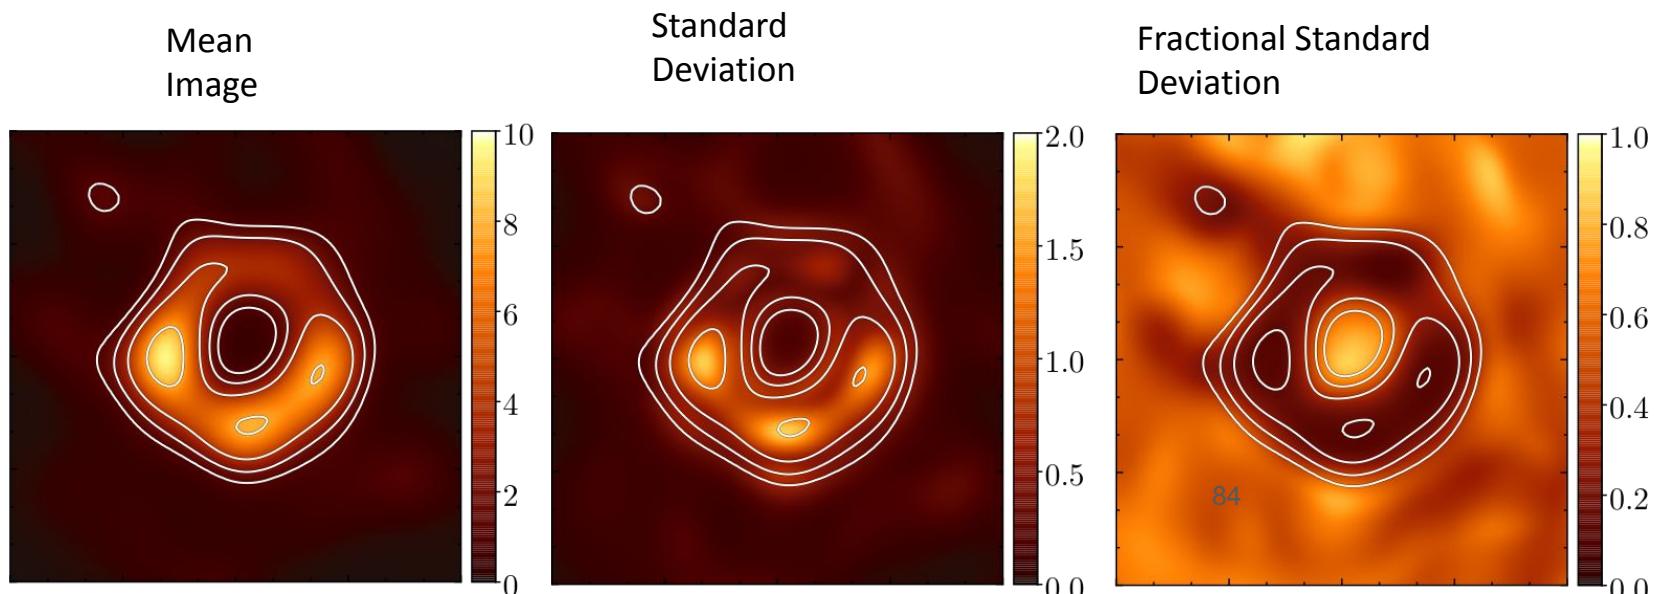
<https://github.com/achael>



Andrew Chael
(Princeton)

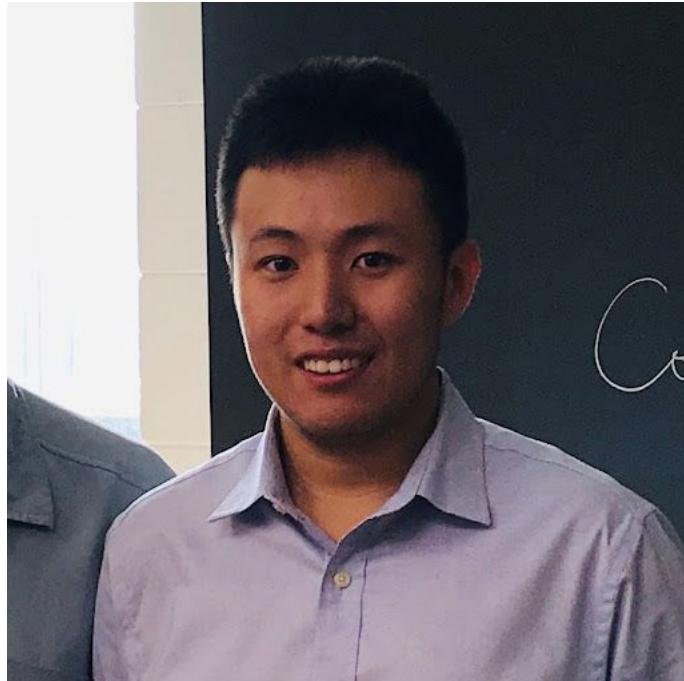


Uncertainty quantification is important!

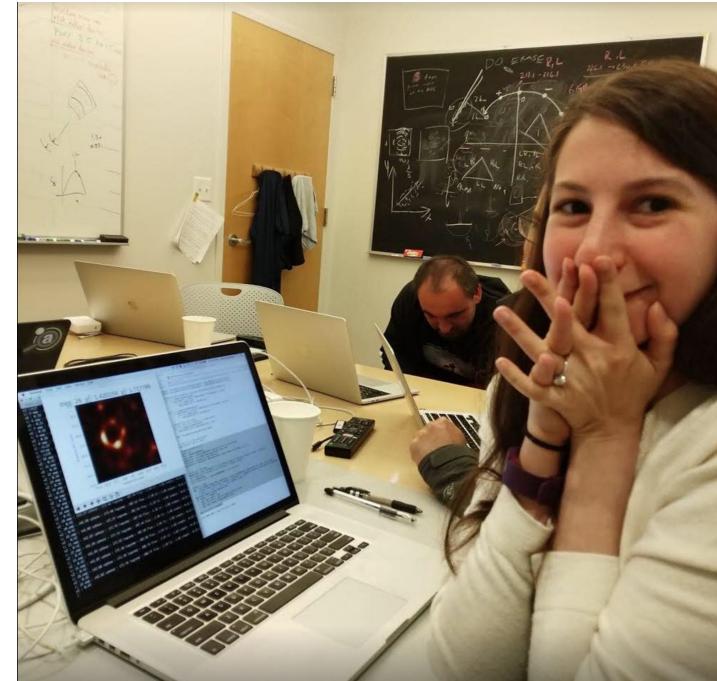


from [M87 Paper IV](#)

Deep Probabilistic Imaging: Uncertainty Quantification and Multi-modal Solution Characterization for Computational Imaging



He Sun (Caltech)



Katie Bouman (Caltech)

Deep Probabilistic Imaging

Big question: Given observable y , what's the distribution of image x ?

THE ASTROPHYSICAL JOURNAL LETTERS, 875:L6 (44pp), 2019 April 10

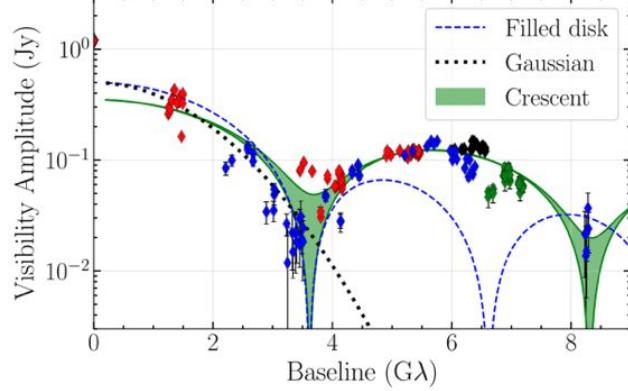
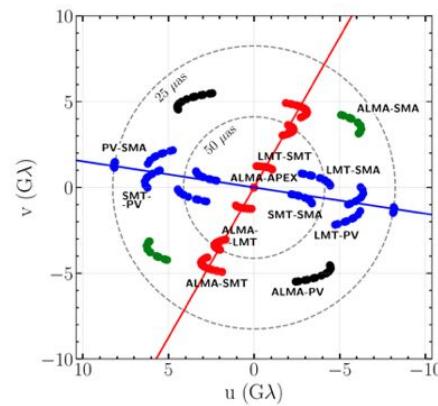
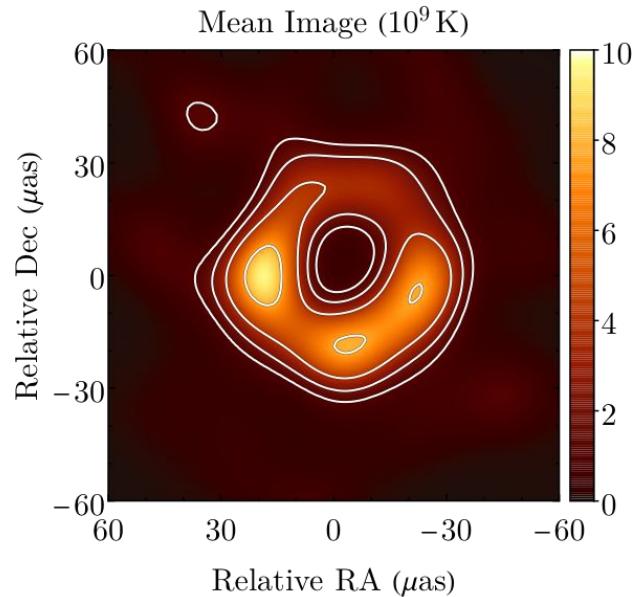


Figure 1. (u, v) -coverage (left panel) and visibility amplitudes (right panel) of M87 for the high-band April 11 data. The (u, v) -coverage has two primary orientations, east-west in blue and north-south in red, with two diagonal filters at large baselines in green and black. Note that the Large Millimeter Telescope (LMT) and the Submillimeter Telescope (SMT) participate in both orientations, and that the LMT amplitudes are subject to significant gain errors. There is evidence for substantial depressions in the visibility amplitudes at ~ 3.4 G λ and ~ 8.3 G λ . The various lines in the right panel show the expected behavior of (dotted line) a Gaussian, (dashed line) a filled disk, and (green area) a crescent shape along different orientations. The image of M87 does not appear to be consistent with a filled disk or a Gaussian.



Regularized Maximum Likelihood (RML)

$$\begin{aligned} \text{Best Image} \\ \hat{\mathbf{x}}_{\text{MAP}} &= \operatorname{argmax}_{\mathbf{x}} [\log p(\mathbf{x}|\mathbf{y})] \\ &= \operatorname{argmax}_{\mathbf{x}} [\log p(\mathbf{y}|\mathbf{x}) + \log p(\mathbf{x})] \\ &\quad \text{Likelihood} \quad \text{Prior} \end{aligned}$$

Posterior Exploration

$$x \sim p(\mathbf{x}|\mathbf{y})$$

The posterior can be sampled directly using Markov Chain Monte Carlo (MCMC) method.

["Hybrid Very Long Baseline Interferometry Imaging and Modeling with themis"](#), ApJ 2020

Posterior Exploration

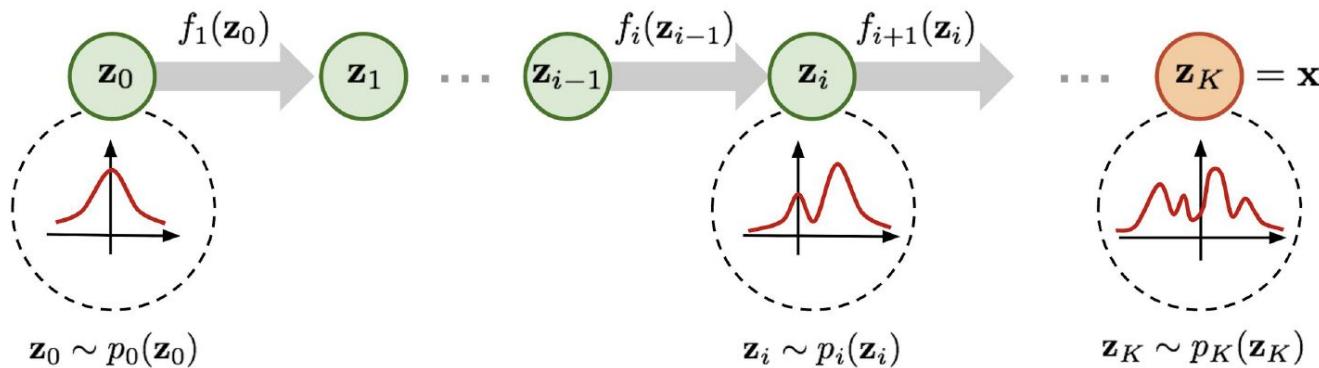
$$x \sim p(\mathbf{x}|\mathbf{y})$$

Or we could use a **Neural Networks** (Normalizing Flow)!

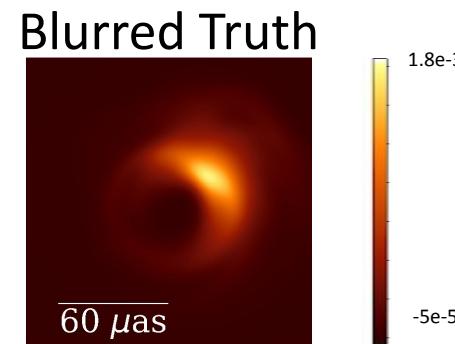
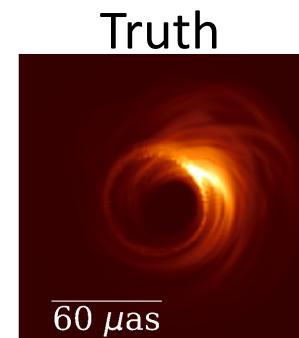
Introduction to Normalizing Flow

Normalizing Flows

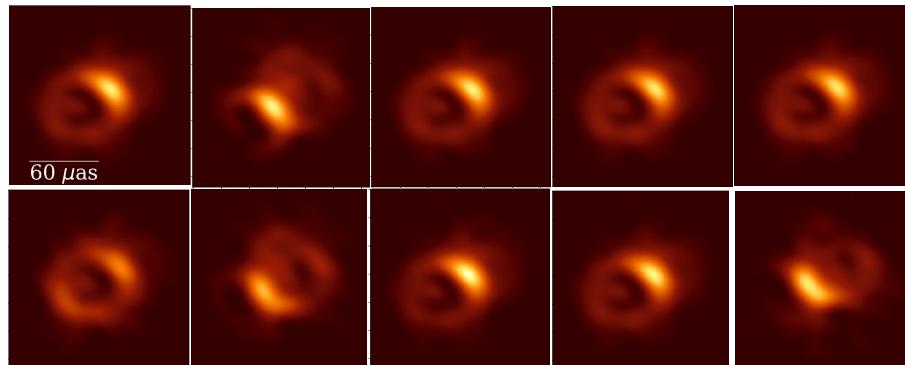
A normalizing flow transforms a simple distribution into a complex one by applying a sequence of invertible transformation functions.



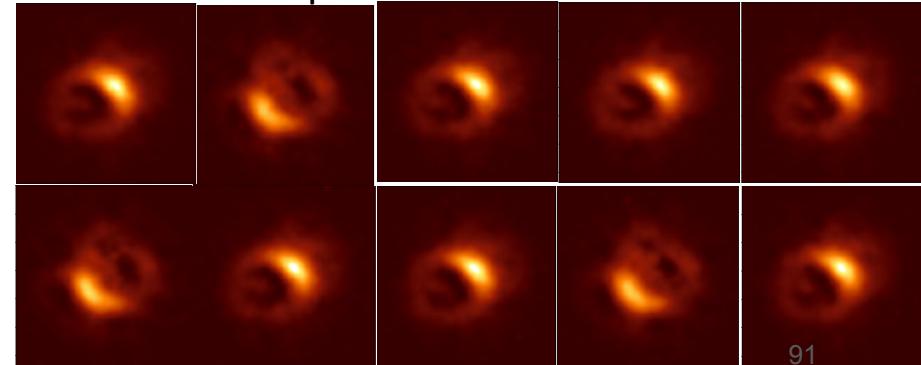
Characterizing Uncertainty & Multimodal Solutions



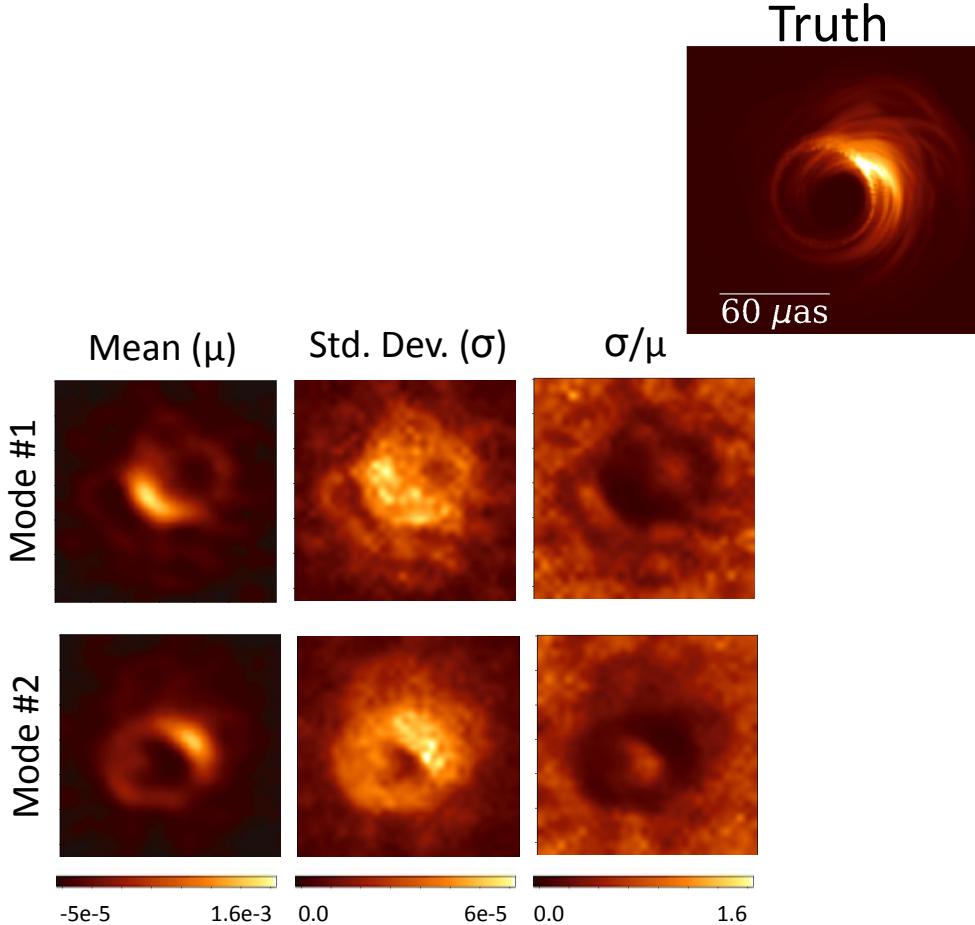
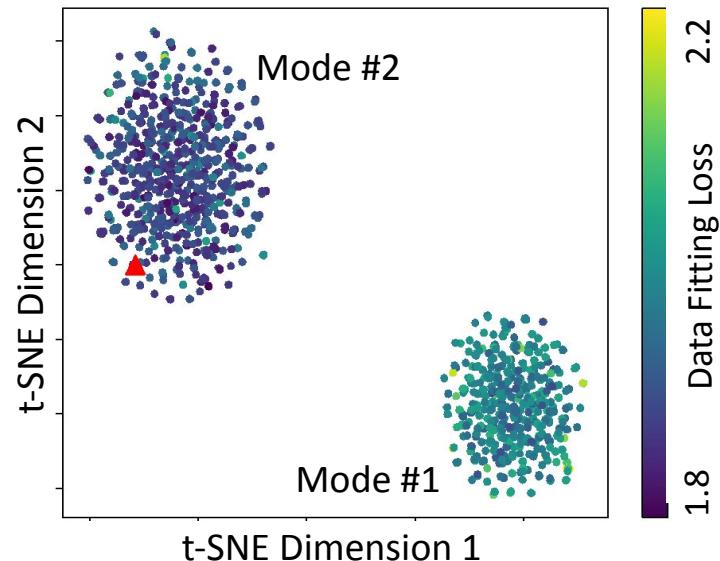
RML results from different initializations



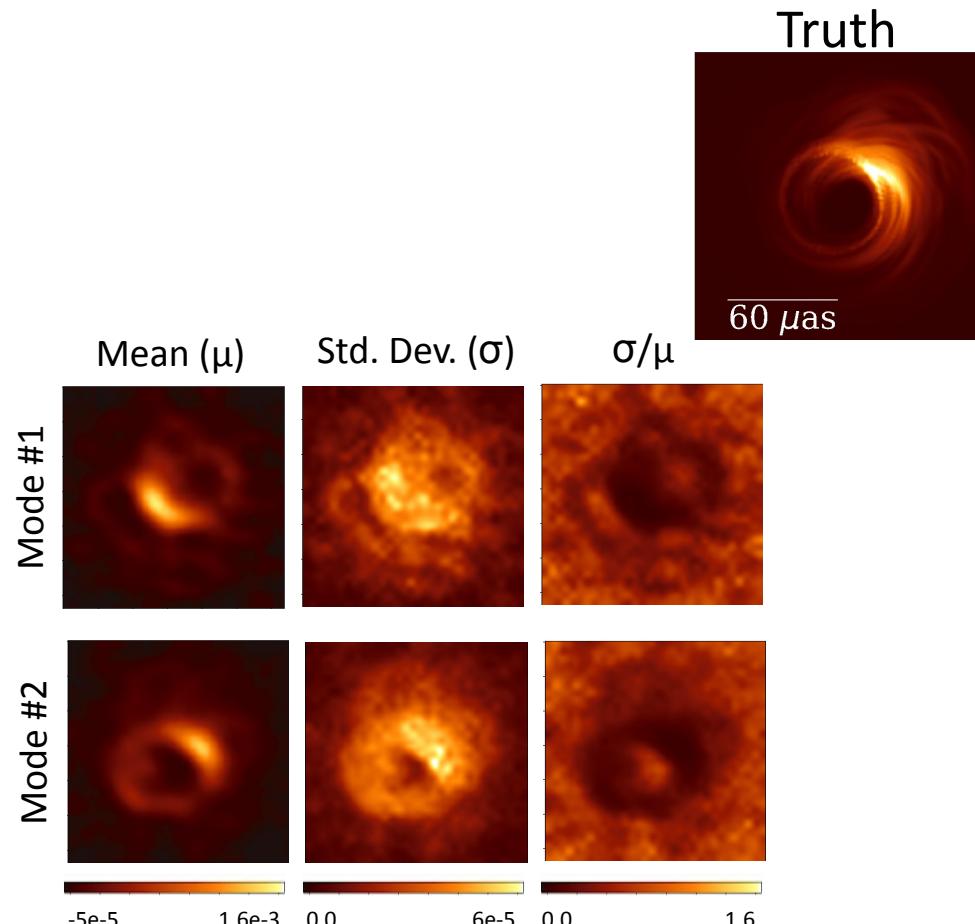
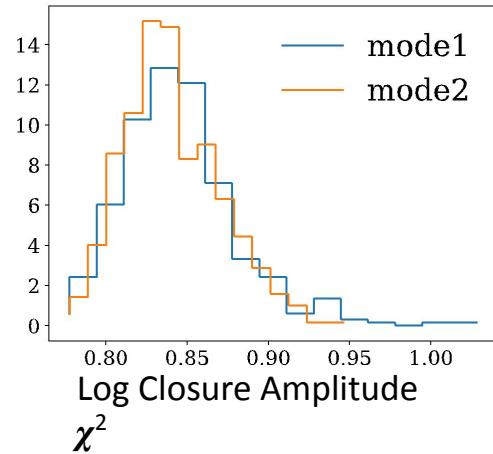
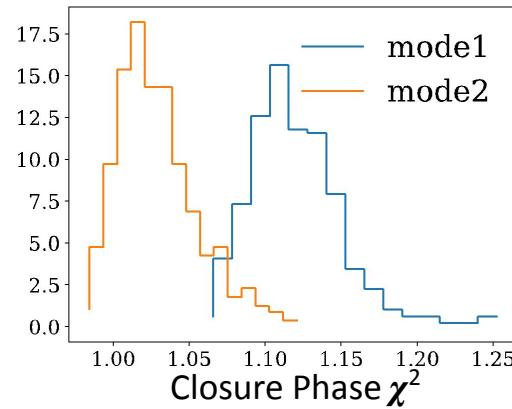
DPI samples from learned network



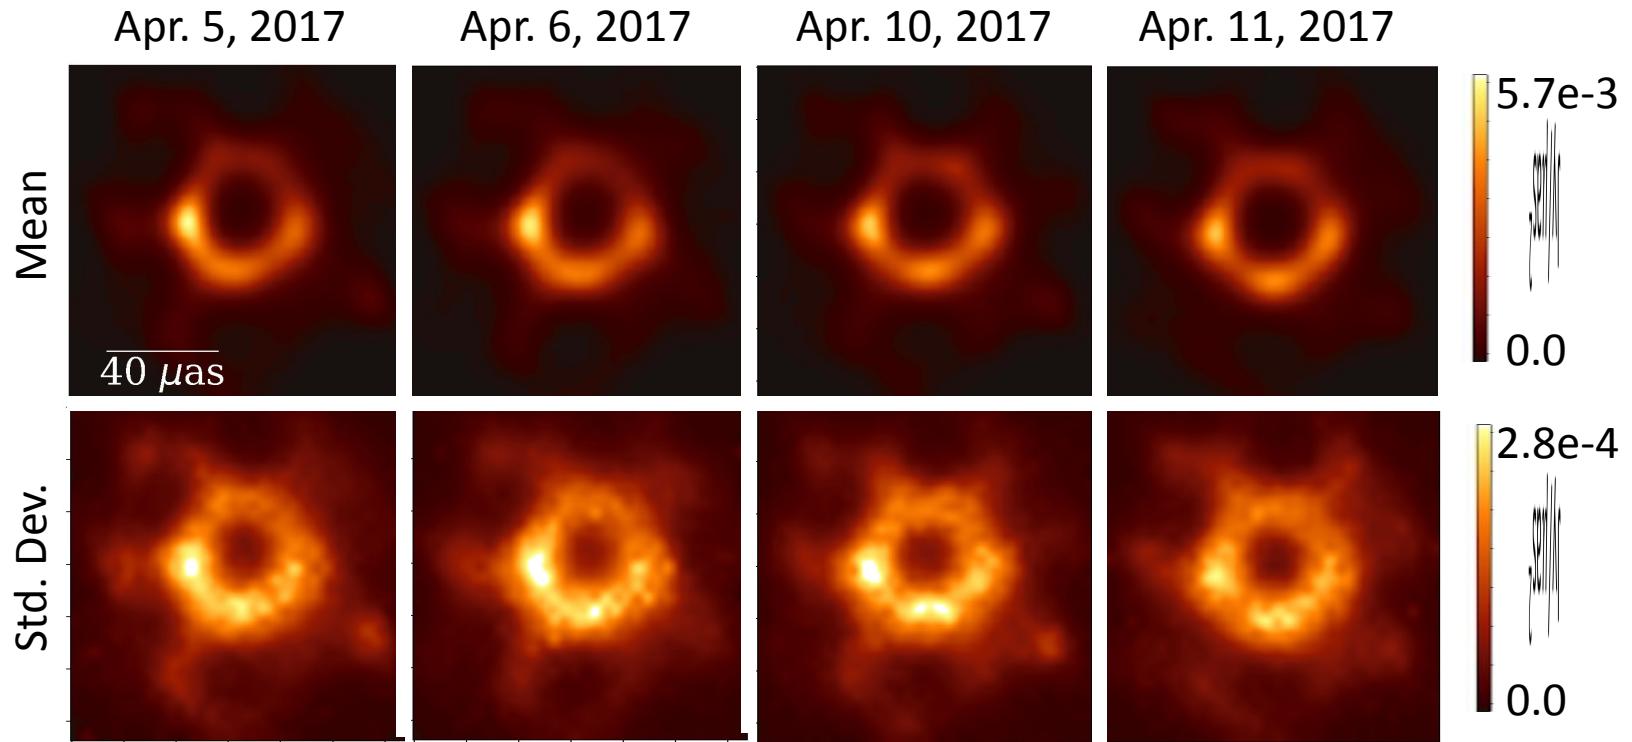
Exploring the Posterior



Exploring the Posterior



Real M87 results: closure quantities + vis amp + MEM+TV2



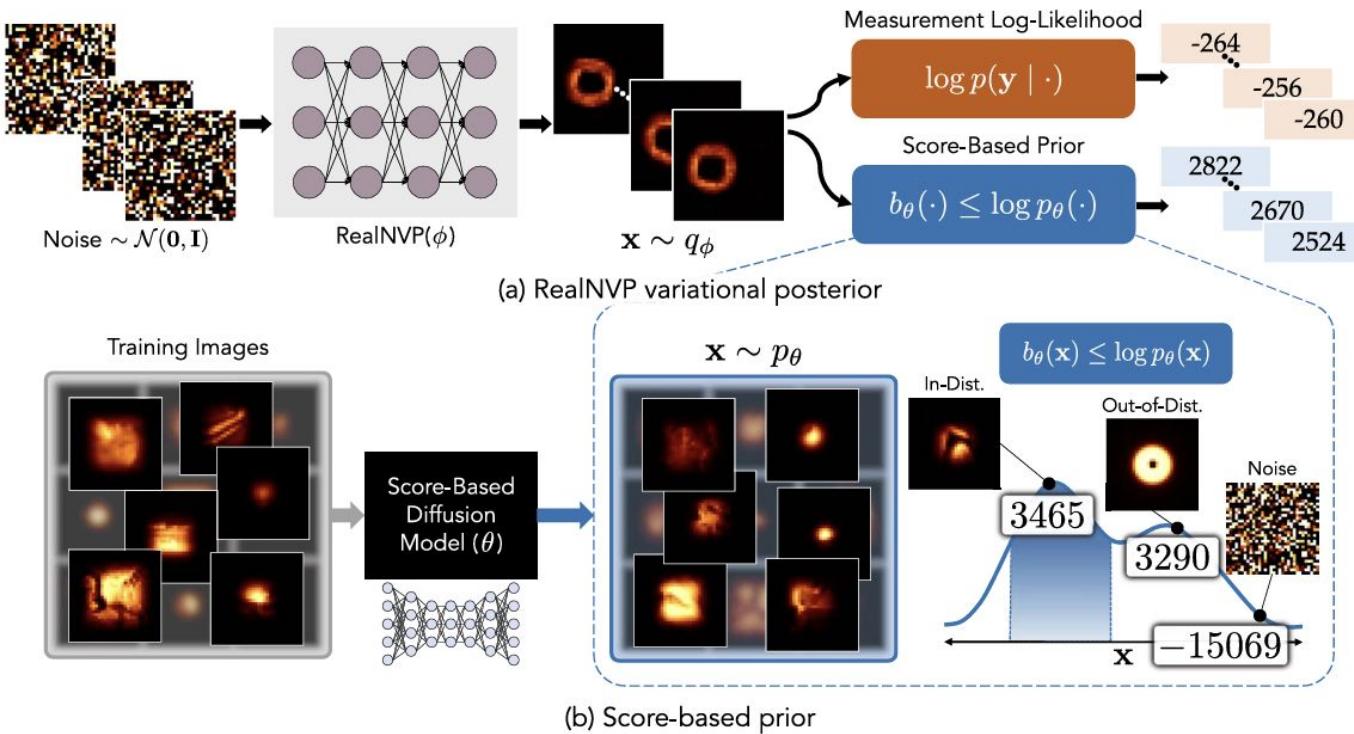
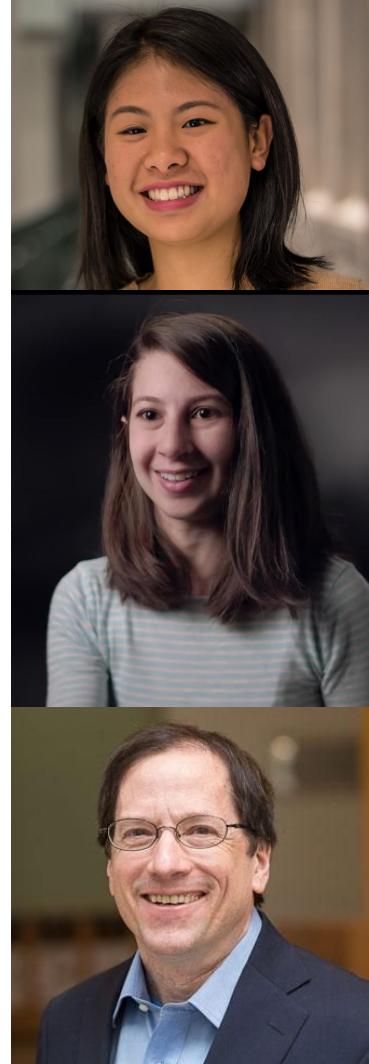
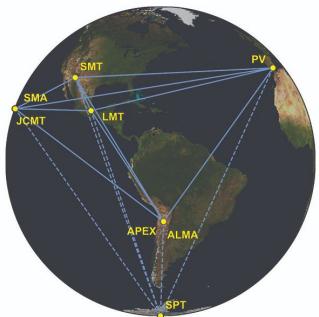
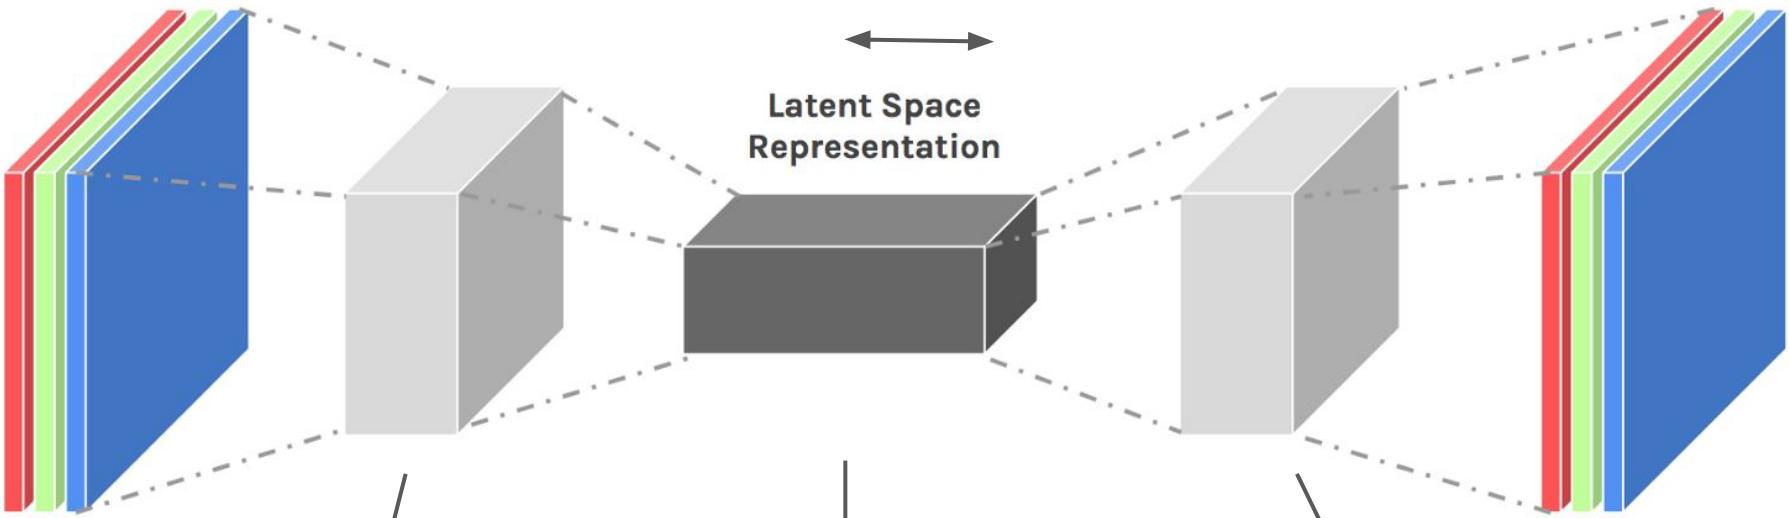
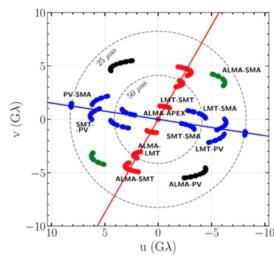


Figure 1. Method illustration. The CIFAR-10 prior was used for these examples; images are shown as 32×32 pixels on a $[0, 1]$ scale. At a high level, we optimize a variational distribution q_ϕ to approximate the image posterior $p_\theta(\cdot | \mathbf{y})$ given a score-based prior p_θ and log likelihood based on EHT measurements. Panel (a) illustrates our particular variational distribution: a RealNVP with parameters ϕ . At each optimization iteration i , the measurement log likelihood (Equation (2)) and the log density under the score-based prior of each sample \mathbf{x} from $q_\phi = q_\phi^{(i)}$ are evaluated. The average gradient is computed with respect to ϕ to update $\phi^{(i)}$. In other words, q_ϕ is optimized to produce samples that have high probability under both measurement likelihood and prior. Panel (b) zooms in to the score-based prior. A score-based prior is based on a score-based diffusion model, a deep generative model with parameters θ , that is trained on images from a target prior. Once trained, the diffusion model generates samples from a generative image distribution p_θ . There is an analytical formula for computing the ELBO $b_\theta(\mathbf{x})$ of the log probability $\log p_\theta(\mathbf{x})$ for any image \mathbf{x} , even for out-of-distribution images and images of pure noise.





Radio Telescope config.



Visibility

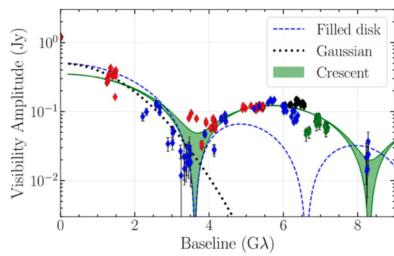
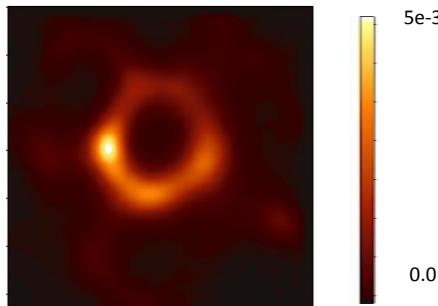


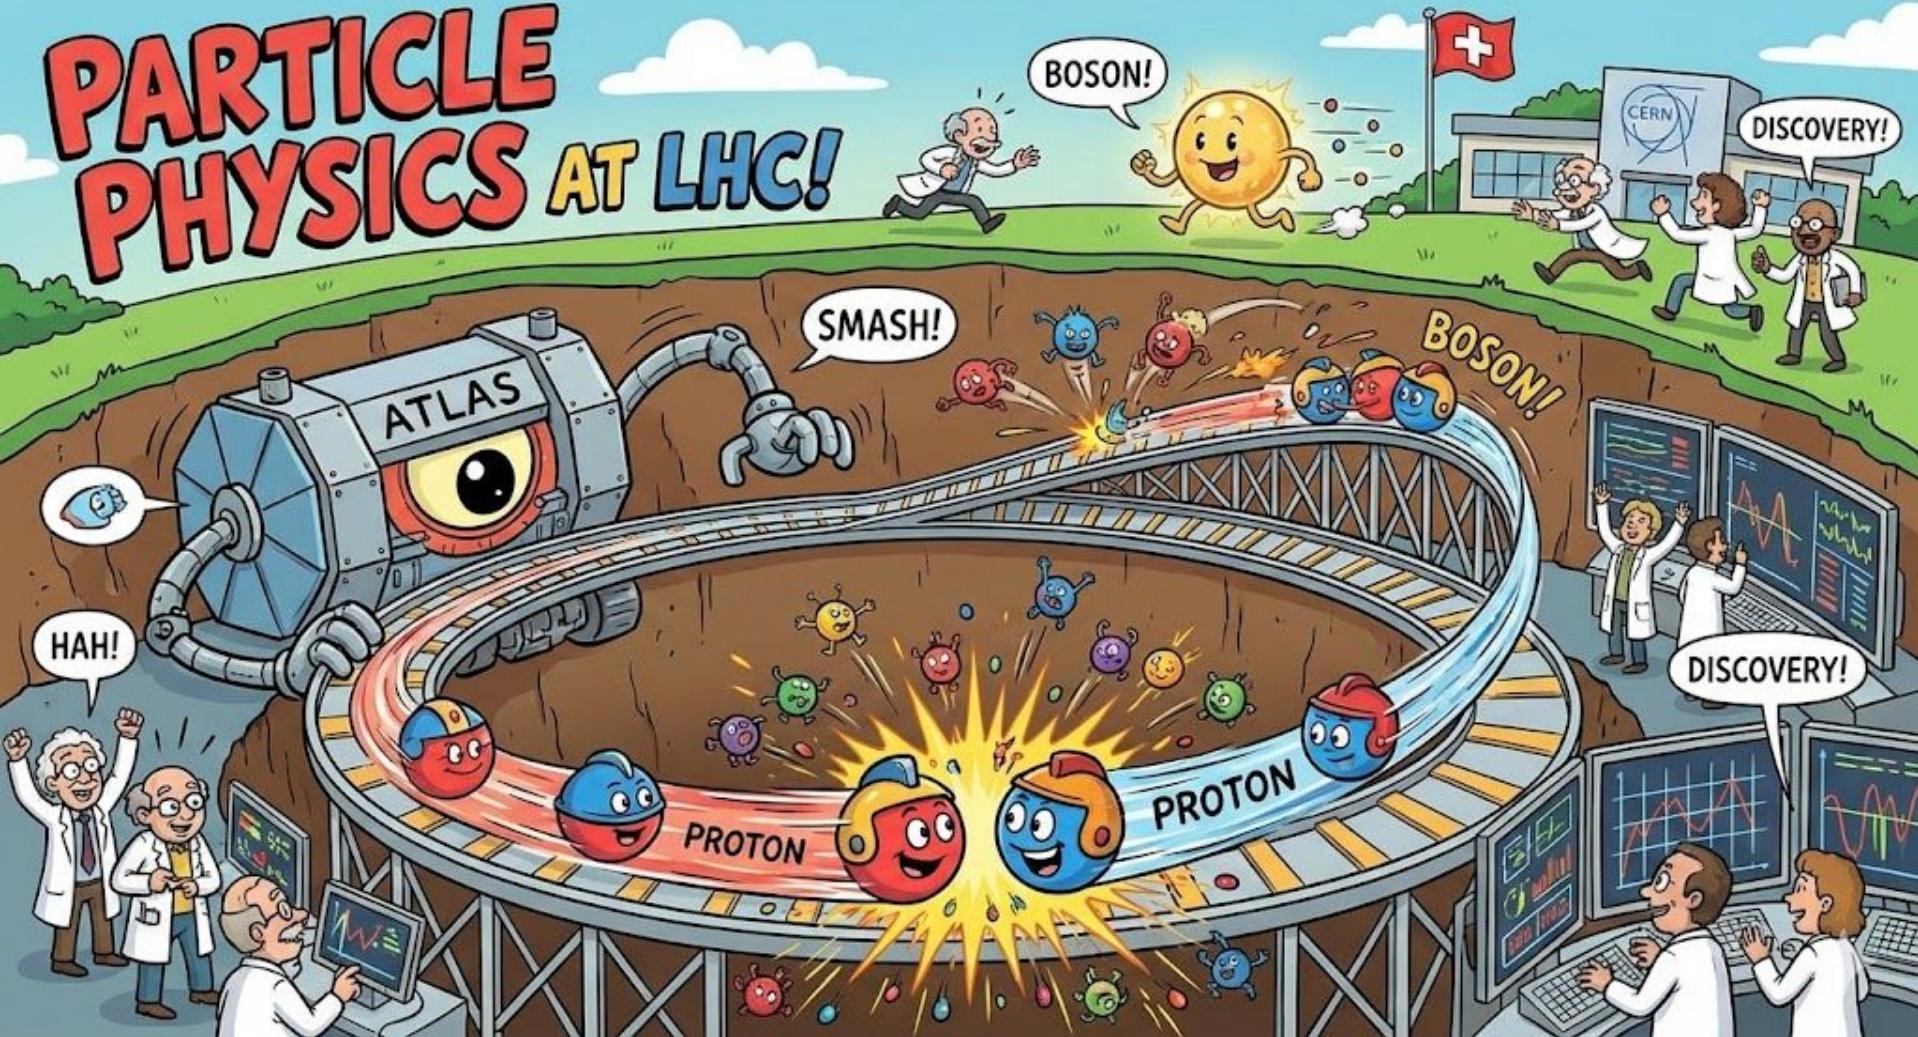
Image Reconstruction Algo. (e.g. RML, CLEAN)

Summary

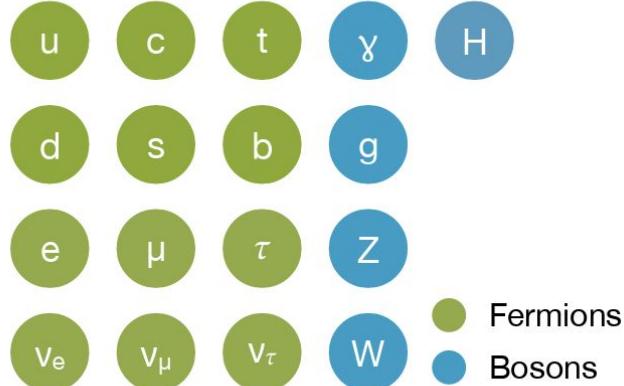
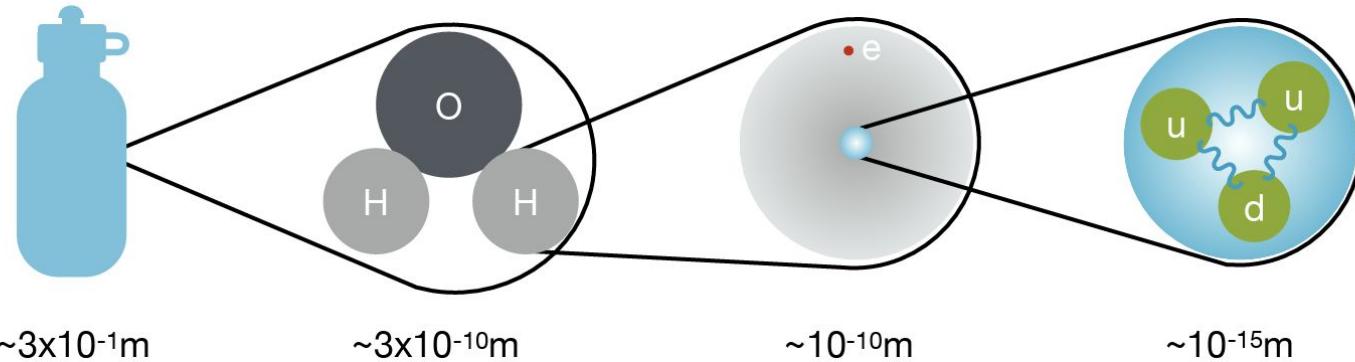
- Deep Probabilistic Imaging (DPI) using variational methods and generative model to approximate the posterior distribution of reconstructed image;
- DPI can capture multiple feasible solutions and quantify the uncertainty;
- DPI is preliminarily tested on EHT simulated data and EHT2017 M87 data.
- Generalizable to other parameter estimation problems



PARTICLE PHYSICS AT LHC!

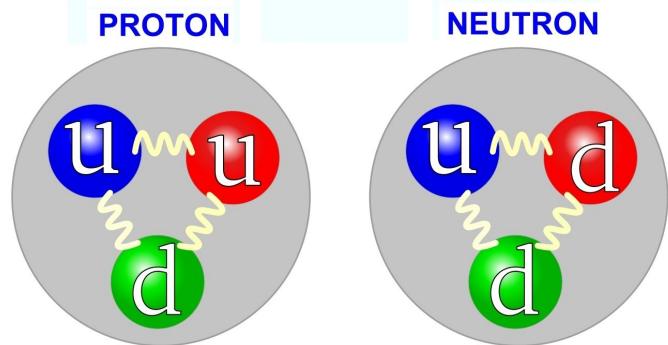


Particle Physics in a Nutshell 1/2



$$\begin{aligned}
 & -\frac{1}{2}\partial_\nu g_\mu^a \partial_\nu g_\mu^a - g_s f^{abc} \partial_\mu g_\nu^a g_\mu^b g_\nu^c - \frac{1}{4}g_s^2 f^{abc} f^{ade} g_\mu^b g_\nu^c g_\mu^d g_\nu^e + \\
 & \frac{1}{2}ig_s^2 (\bar{q}_i^\sigma \gamma^\mu q_j^\sigma) g_\mu^a + G^a \partial^2 G^a + g_s f^{abc} \partial_\mu G^a G^b g_\mu^c - \partial_\nu W_\mu^+ \partial_\nu W_\mu^- - \\
 & M^2 W_\mu^+ W_\mu^- - \frac{1}{2}\partial_\nu Z_\mu^0 \partial_\nu Z_\mu^0 - \frac{1}{2c_w^2} M^2 Z_\mu^0 Z_\mu^0 - \frac{1}{2}\partial_\mu A_\nu \partial_\mu A_\nu - \frac{1}{2}\partial_\mu H \partial_\mu H - \\
 & \frac{1}{2}m_h^2 H^2 - \partial_\mu \phi^+ \partial_\mu \phi^- - M^2 \phi^+ \phi^- - \frac{1}{2}\partial_\mu \phi^0 \partial_\mu \phi^0 - \frac{1}{2c_w^2} M \phi^0 \phi^0 - \beta_h \left[\frac{2M^2}{g^2} + \right. \\
 & \left. \frac{2M}{g} H + \frac{1}{2}(H^2 + \phi^0 \phi^0 + 2\phi^+ \phi^-) \right] + \frac{2M^4}{g^2} \alpha_h - ig c_w [\partial_\nu Z_\mu^0 (W_\mu^+ W_\nu^- - \\
 & W_\nu^+ W_\mu^-) - Z_\nu^0 (W_\mu^+ \partial_\nu W_\mu^- - W_\mu^- \partial_\nu W_\mu^+) + Z_\mu^0 (W_\nu^+ \partial_\nu W_\mu^- - \\
 & W_\nu^- \partial_\nu W_\mu^+)] - ig s_w [\partial_\nu A_\mu (W_\mu^+ W_\nu^- - W_\nu^+ W_\mu^-) - A_\nu (W_\mu^+ \partial_\nu W_\mu^- - \\
 & W_\mu^- \partial_\nu W_\mu^+) + A_\mu (W_\nu^+ \partial_\nu W_\mu^- - W_\nu^- \partial_\nu W_\mu^+)] - \frac{1}{2}g^2 W_\mu^+ W_\mu^- W_\nu^+ W_\nu^- + \\
 & \frac{1}{2}g^2 W_\mu^+ W_\nu^- W_\mu^+ W_\nu^- + g^2 c_w^2 (Z_\mu^0 W_\mu^+ Z_\nu^0 W_\nu^- - Z_\mu^0 Z_\mu^0 W_\nu^+ W_\nu^-) + \\
 & g^2 s_w^2 (A_\mu W_\mu^+ A_\nu W_\nu^- - A_\mu A_\nu W_\mu^+ W_\nu^-) + g^2 s_w c_w [A_\mu Z_\nu^0 (W_\mu^+ W_\nu^- - \\
 & W_\nu^+ W_\mu^-) - 2A_\mu Z_\nu^0 W_\nu^+ W_\nu^-] - g\alpha [H^3 + H\phi^0 \phi^0 + 2H\phi^+ \phi^-] - \\
 & \frac{1}{8}g^2 \alpha_h [H^4 + (\phi^0)^4 + 4(\phi^+ \phi^-)^2 + 4(\phi^0)^2 \phi^+ \phi^- + 4H^2 \phi^+ \phi^- + 2(\phi^0)^2 H^2] - \\
 & \dots
 \end{aligned}$$

Standard Model of Elementary Particles

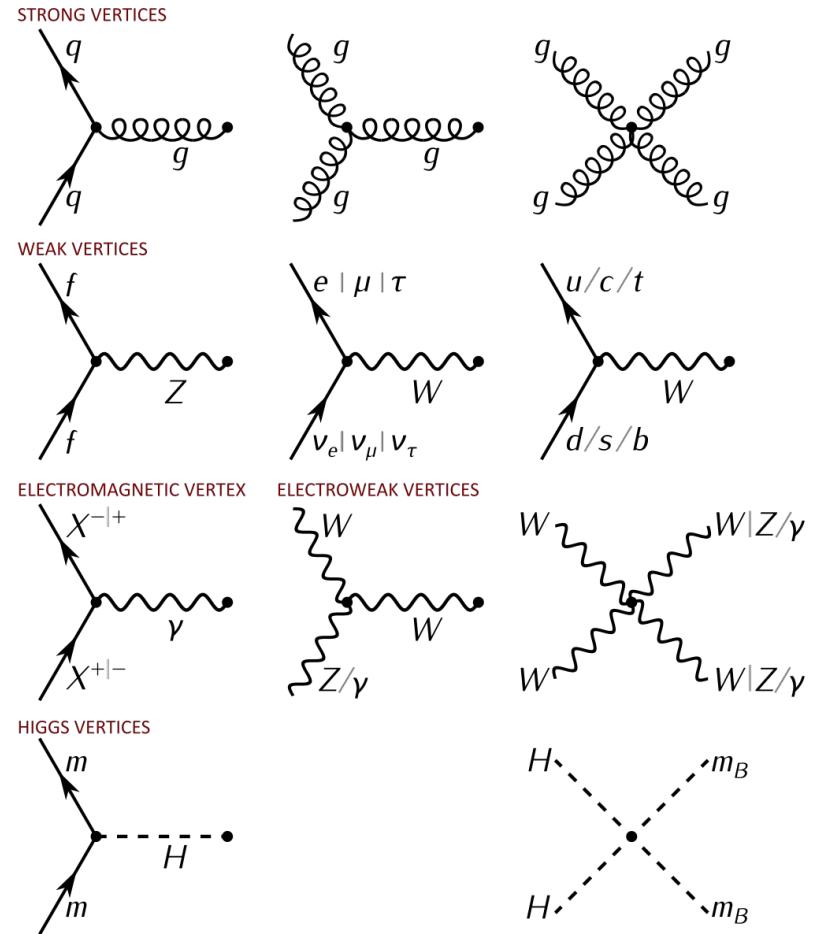


three generations of matter (fermions)			interactions / force carriers (bosons)	
QUARKS	I	II	III	
	mass $=2.16 \text{ MeV}/c^2$	mass $=1.273 \text{ GeV}/c^2$	mass $=172.57 \text{ GeV}/c^2$	0
	charge $\frac{2}{3}$	charge $\frac{2}{3}$	charge $\frac{2}{3}$	0
	spin $\frac{1}{2}$	spin $\frac{1}{2}$	spin $\frac{1}{2}$	1
	u	c	t	g
	up	charm	top	gluon
	d	s	b	γ
	down	strange	bottom	photon
				Z boson
LEPTONS	$\approx 0.511 \text{ MeV}/c^2$	$\approx 105.66 \text{ MeV}/c^2$	$\approx 1.77693 \text{ GeV}/c^2$	$\approx 91.188 \text{ GeV}/c^2$
	-1 $\frac{1}{2}$	-1 $\frac{1}{2}$	-1 $\frac{1}{2}$	0 1
	e	μ	τ	Z boson
	electron	muon	tau	W boson
GAUGE BOSONS VECTOR BOSONS	$<0.8 \text{ eV}/c^2$	$<0.17 \text{ MeV}/c^2$	$<18.2 \text{ MeV}/c^2$	$\approx 80.3692 \text{ GeV}/c^2$
	0 $\frac{1}{2}$	0 $\frac{1}{2}$	0 $\frac{1}{2}$	± 1 1
	ν_e electron neutrino	ν_μ muon neutrino	ν_τ tau neutrino	W boson
				SCALAR BOSONS

Feynman Diagram



Richard Feynman

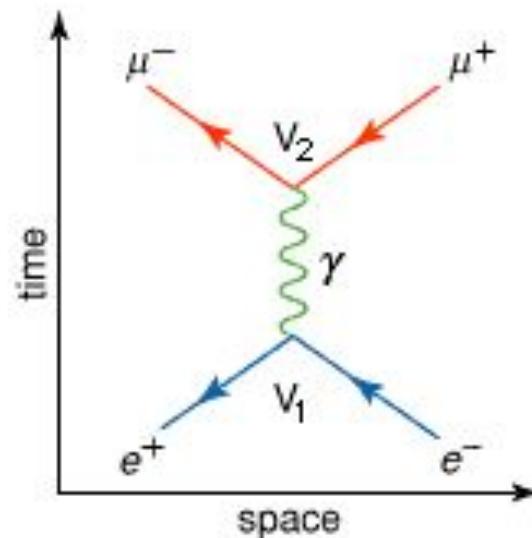


Feynman Diagram 101



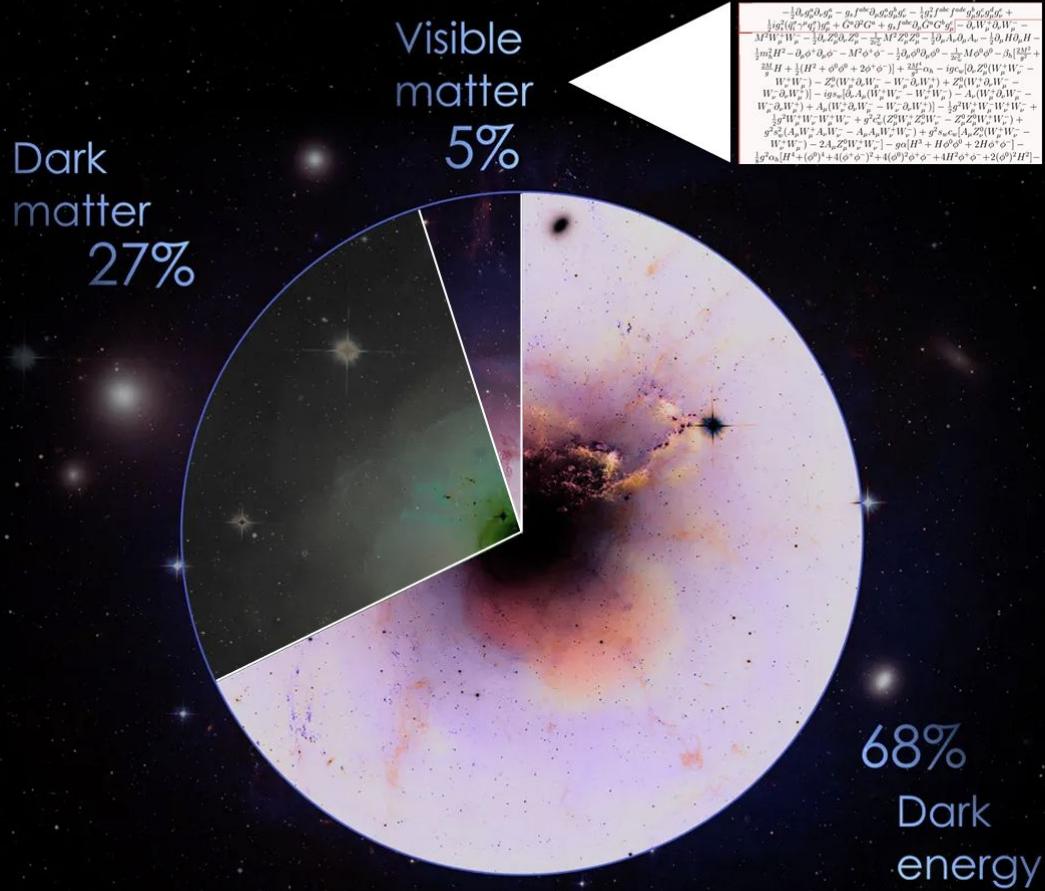
Richard Feynman

$$e^+ + e^- \rightarrow \mu^- + \mu^+$$

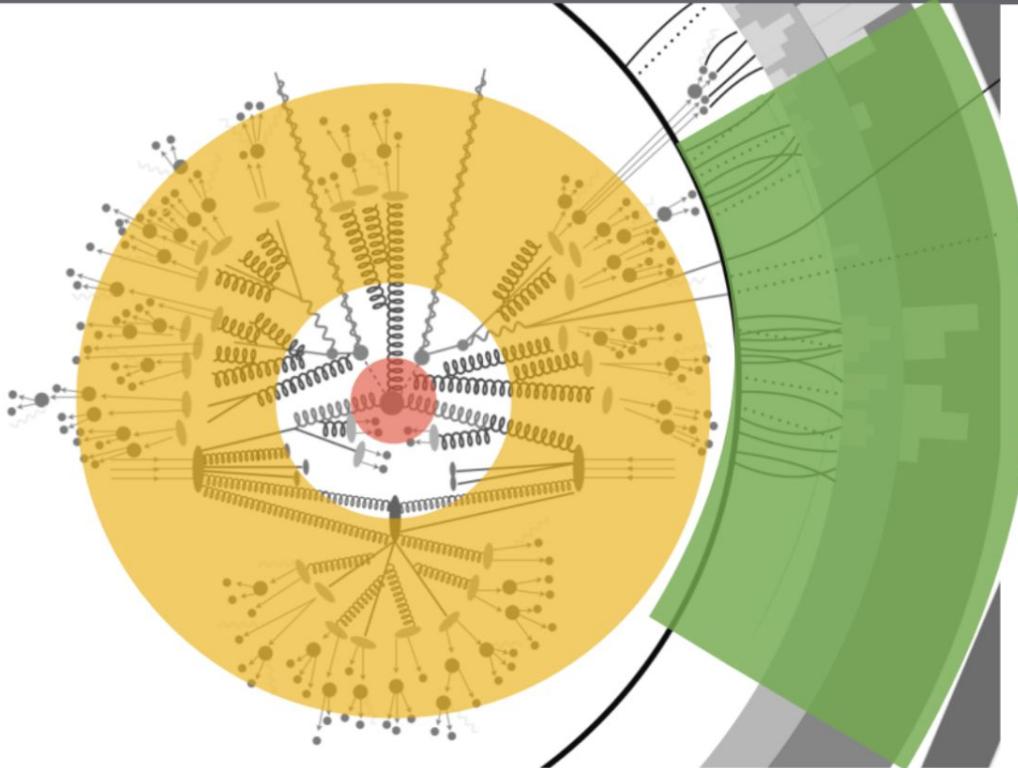


© 2007 Encyclopædia Britannica, Inc.

Particle Physics in a Nutshell 2/2



Very Precise Simulations Available ... But Expensive!



Hard Process

Showering &

Detector

Hadronization

Interaction

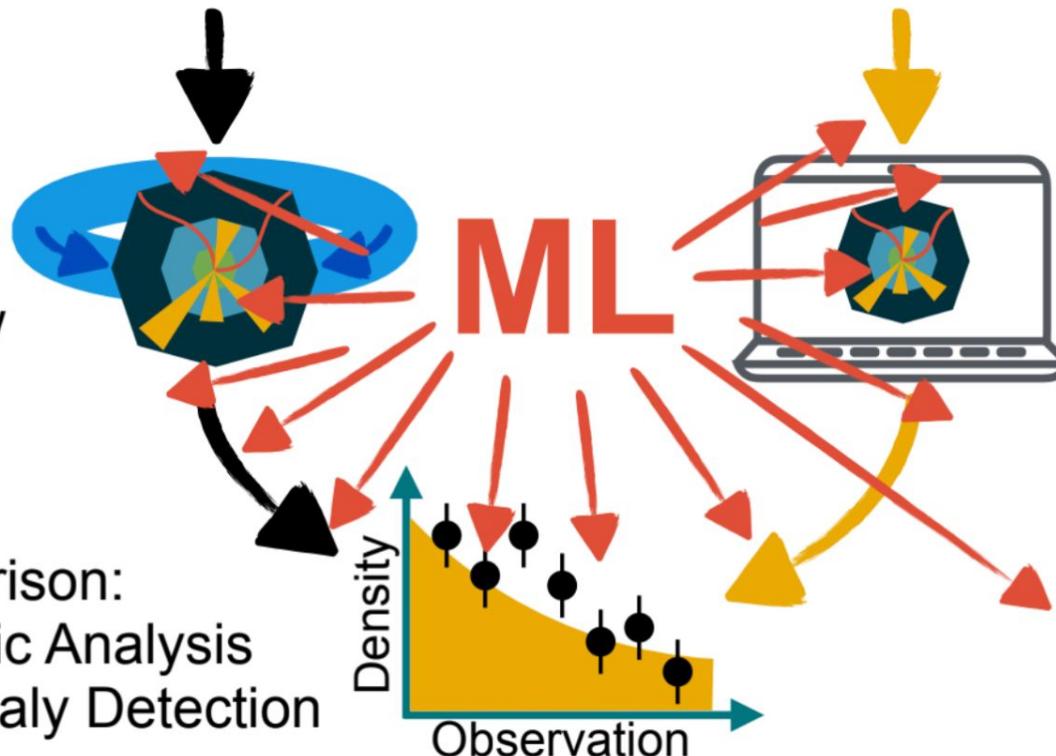
$$\begin{aligned} & -\frac{1}{2} \partial_\mu g_{\mu\nu}^a \partial_\nu g_{\mu\nu}^a - g_\mu f^{\mu\nu} \partial_\nu g_{\mu\nu}^a g_\nu^a - \frac{1}{2} g_{\mu\nu}^2 f^{\mu\nu} f^{\mu\nu} g_{\mu\nu}^a g_\nu^a + \\ & + \frac{1}{2} g_{\mu\nu}^2 g_{\mu\nu}^a g_{\mu\nu}^a + G^\mu \partial_\mu G^\nu + g_\mu g_{\mu\nu}^a g_\nu^a G^\mu G^\nu - g_\mu^2 W_\mu^a W_\mu^a - \\ & - M^2 W_\mu^a W_\mu^a - \frac{1}{2} g_{\mu\nu}^2 g_{\mu\nu}^a g_{\mu\nu}^a W_\mu^a W_\mu^a - \frac{1}{2} g_{\mu\nu}^2 g_{\mu\nu}^a g_\nu^a H_\mu^a H_\nu^a + \\ & + \frac{1}{2} g_{\mu\nu}^2 H_\mu^a H_\nu^a - \partial_\mu \phi^\dagger \partial_\nu \phi - M^2 \phi^\dagger \phi - \frac{1}{2} g_{\mu\nu}^2 \partial_\mu \phi^\dagger \partial_\nu \phi - \frac{1}{2} g_{\mu\nu}^2 M^2 \phi^\dagger \phi - g_\mu^2 \frac{128}{3} + \\ & + \frac{24}{5} H + \frac{1}{2} H^2 - 4 \phi^\dagger \phi + 4 \phi^\dagger \phi - \frac{128}{3} M^2 \phi^\dagger \phi - 16 g_\mu^2 (4 Z_2^0 W_\mu^a W_\mu^a - \\ & - W_\mu^a W_\mu^a) - 2 g_\mu^2 (4 Z_3^0 W_\mu^a W_\mu^a - W_\mu^a W_\mu^a) + 2 Z_2^0 (4 Z_3^0 W_\mu^a W_\mu^a - \\ & - W_\mu^a W_\mu^a) - 16 g_\mu^2 (4 Z_3^0 W_\mu^a W_\mu^a - W_\mu^a W_\mu^a) - 4 Z_3^0 (4 Z_3^0 W_\mu^a W_\mu^a - \\ & - W_\mu^a W_\mu^a) + 4 (W_\mu^a W_\mu^a - W_\mu^a W_\mu^a) - \frac{1}{2} g_\mu^2 (W_\mu^a W_\mu^a W_\mu^a W_\mu^a + \\ & + 4 g_\mu^2 W_\mu^a W_\mu^a W_\mu^a W_\mu^a + g^2 c_\mu^2 (Z_2^0 W_\mu^a Z_2^0 W_\mu^a - Z_2^0 Z_2^0 W_\mu^a W_\mu^a) + \\ & + g^2 c_\mu^2 (4 Z_3^0 W_\mu^a W_\mu^a - 4 g_\mu^2 W_\mu^a W_\mu^a) + g^2 c_\mu^2 c_\mu (4 Z_3^0 W_\mu^a W_\mu^a - \\ & - W_\mu^a W_\mu^a) - 2 A_\mu Z_2^0 W_\mu^a W_\mu^a) - g_\mu (H^2 + H \phi^\dagger \phi + 2 H \phi^\dagger \phi) - \\ & - g^2 c_\mu (H^4 + (\phi^\dagger)^4 + 4 (\phi^\dagger)^2 \phi^\dagger \phi + 4 (\phi^\dagger)^2 \phi^\dagger \phi + 2 (\phi^\dagger)^2 H^2) - \end{aligned}$$

~100s/event
>99% of time

Nature

Experiment:

- Tracking
- Particle Flow



Theory

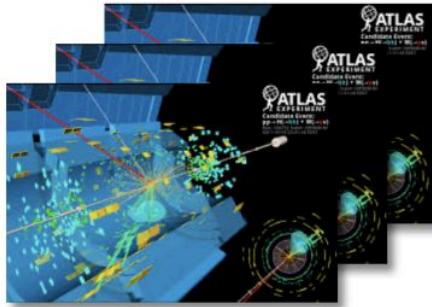
Simulation:

- GANs
- CFMs

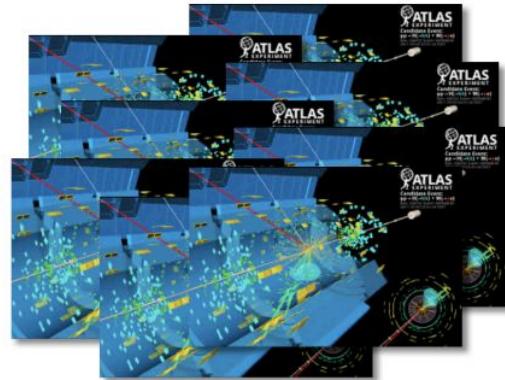
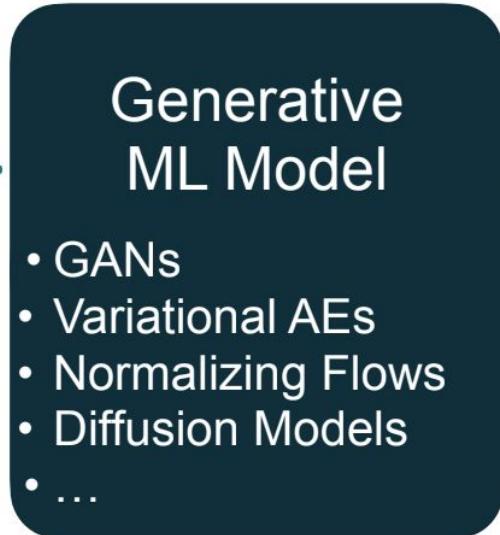
Comparison:

- Classic Analysis
- Anomaly Detection

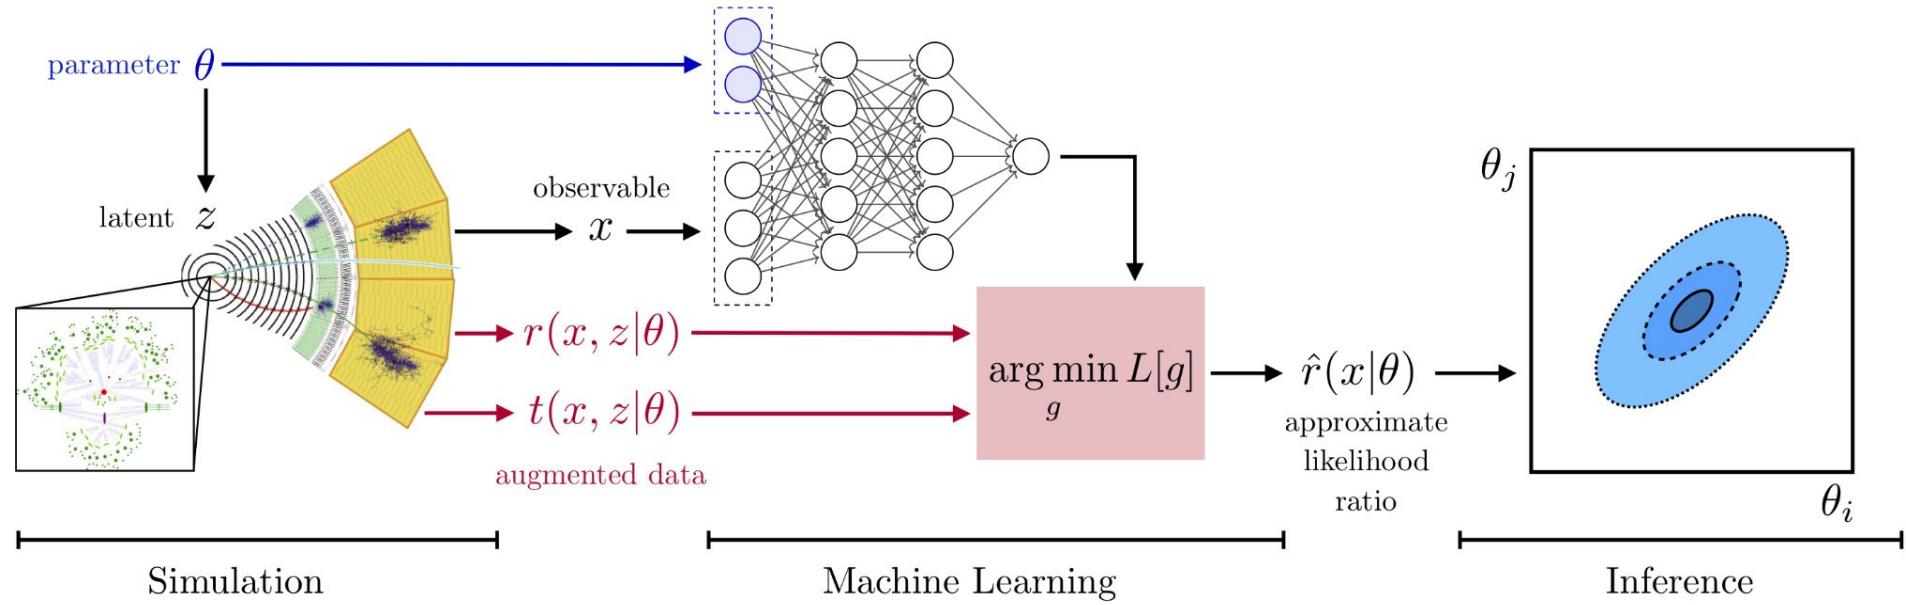
**Foundation
Models**



Simulation /
Recorded Data

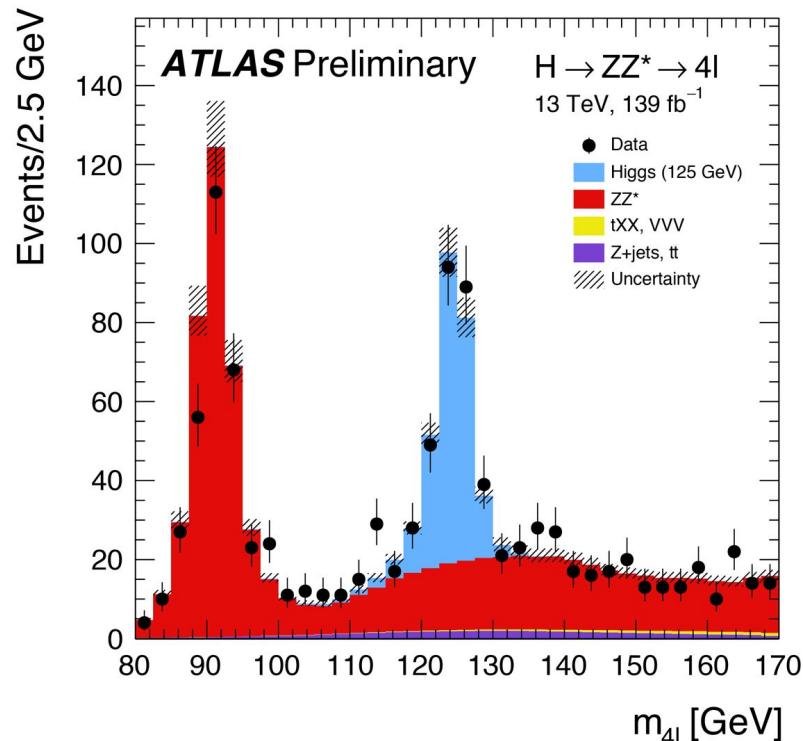
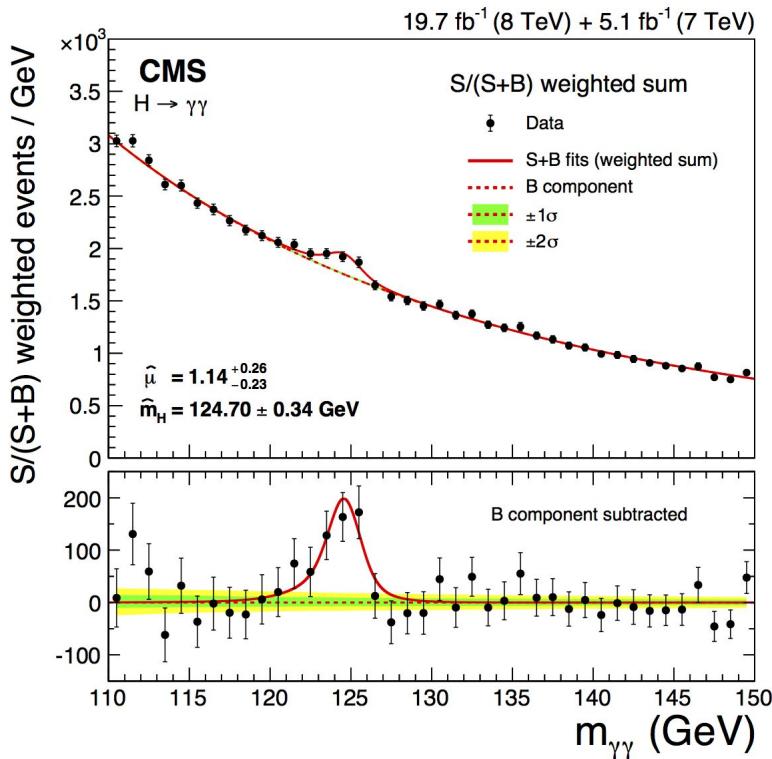


Oversampled /
Interpolated



K. Cranmer, J. Brehmer and G. Louppe

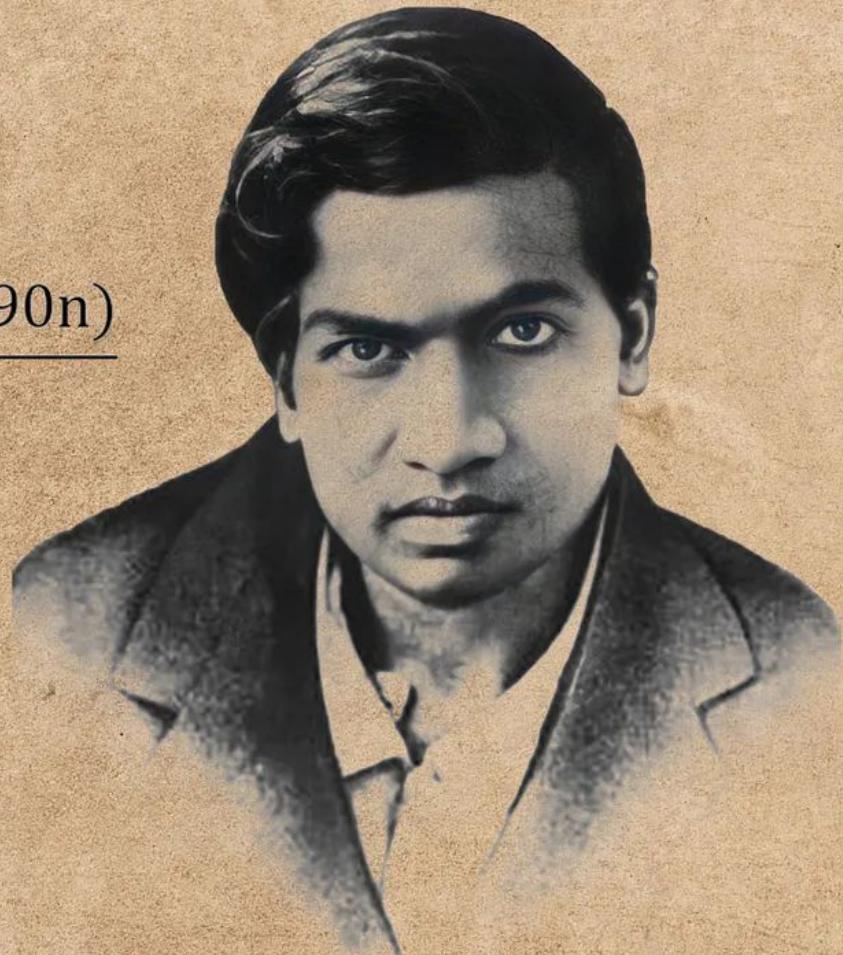
What analysis looks like...



AI for Math

$$\frac{1}{\pi} = \frac{2\sqrt{2}}{9801} \sum_{n=0}^{\infty} \frac{(4n)! (1103 + 26390n)}{(n!)^4 396^{4n}}$$

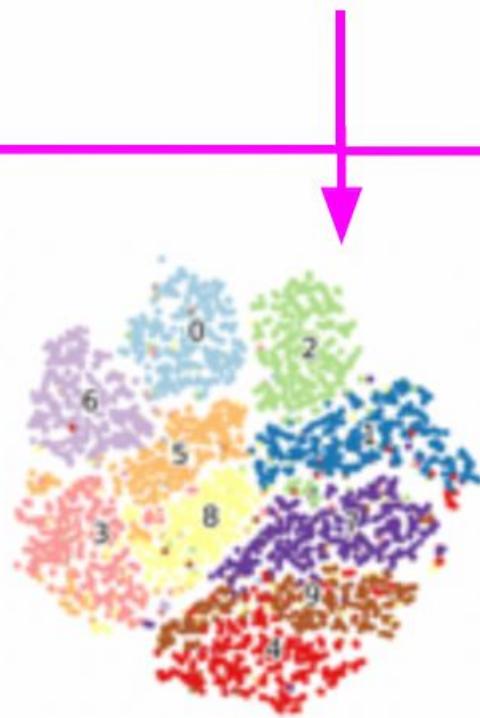
$$\pi = \frac{9801}{1103\sqrt{8}}$$



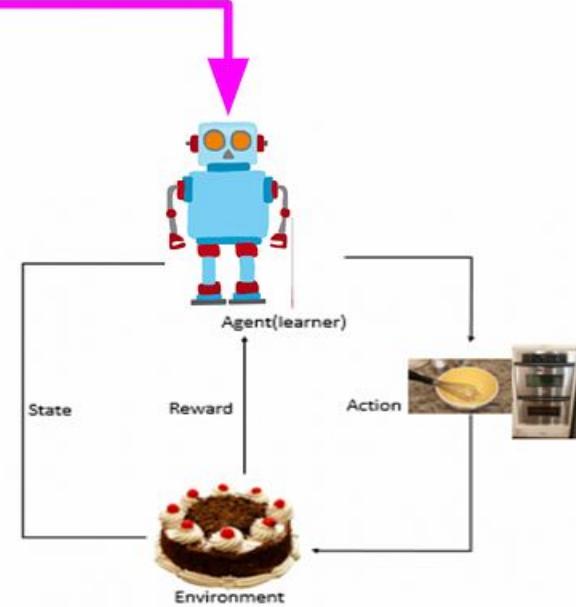
Machine Learning



Supervised Learning

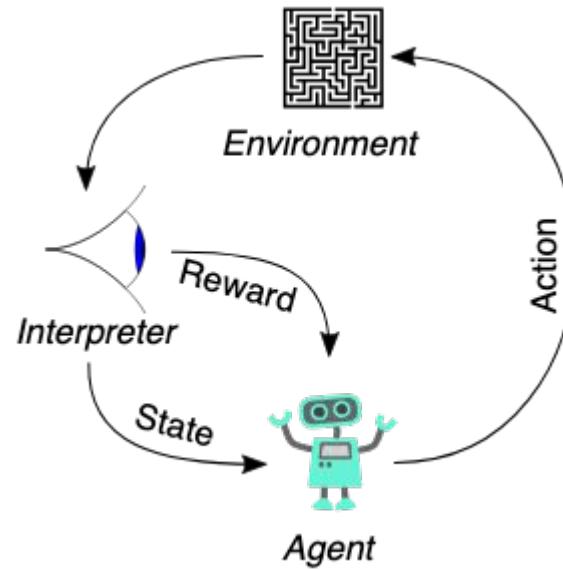


Unsupervised Learning



Reinforcement Learning

What's Reinforcement Learning



Successful example of Reinforcement Learning



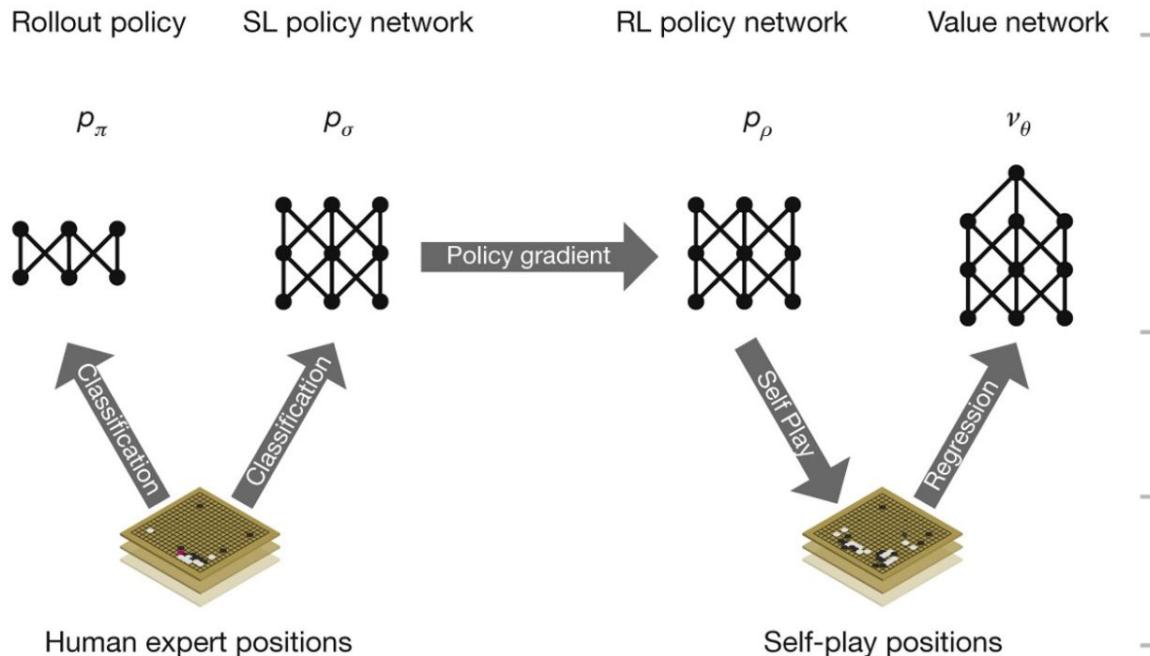
AlphaGo



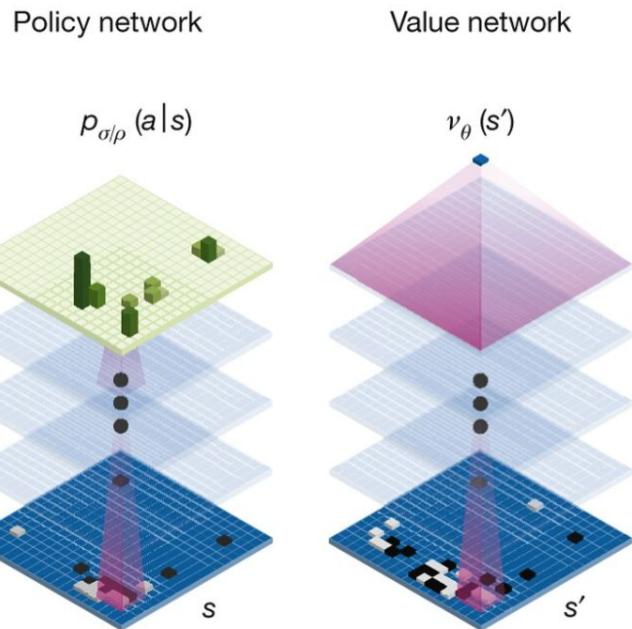
LLM with CoT +
Reinforcement Learning

AlphaGo Fan/Lee (2016)

a



b



Learning from Human Expert

Move



Generating new self-play to prevent overfitting in RL



Reinforcement learning of value networks

The final stage of the training pipeline focuses on position evaluation, estimating a value function $v^p(s)$ that predicts the outcome from position s of games played by using policy p for both players²⁸⁻³⁰

$$v^p(s) = \mathbb{E}[z_t | s_t = s, a_{t \dots T} \sim p]$$

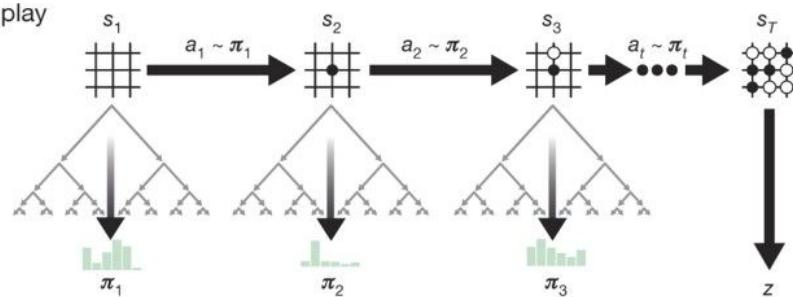
Ideally, we would like to know the optimal value function under perfect play $v^*(s)$; in practice, we instead estimate the value function v^{p_θ} for our strongest policy, using the RL policy network p_θ . We approximate the value function using a value network $v_\theta(s)$ with weights θ , $v_\theta(s) \approx v^{p_\theta}(s) \approx v^*(s)$. This neural network has a similar architecture to the policy network, but outputs a single prediction instead of a probability distribution. We train the weights of the value network by regression on state-outcome pairs (s, z) , using stochastic gradient descent to minimize the mean squared error (MSE) between the predicted value $v_\theta(s)$, and the corresponding outcome z

$$\Delta\theta \propto \frac{\partial v_\theta(s)}{\partial \theta} (z - v_\theta(s))$$

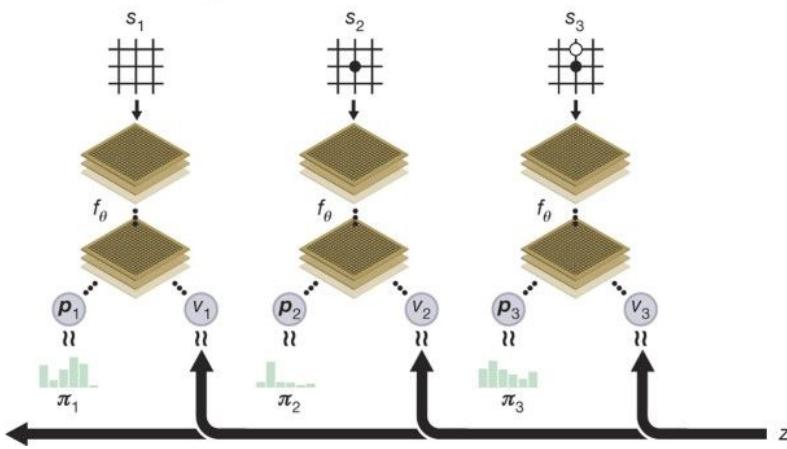
The naive approach of predicting game outcomes from data consisting of complete games leads to overfitting. The problem is that successive positions are strongly correlated, differing by just one stone, but the regression target is shared for the entire game. When trained on the KGS data set in this way, the value network memorized the game outcomes rather than generalizing to new positions, achieving a minimum MSE of 0.37 on the test set, compared to 0.19 on the training set. To mitigate this problem, we generated a new self-play data set consisting of 30 million distinct positions, each sampled from a separate game. Each game was played between the RL policy network and itself until the game terminated. Training on this data set led to MSEs of 0.226 and 0.234 on the training and test set respectively, indicating minimal overfitting. Figure 2b shows the position evaluation accuracy of the value network, compared to Monte Carlo rollouts using the fast rollout policy p_π ; the value function was consistently more accurate. A single evaluation of $v_\theta(s)$ also approached the accuracy of Monte Carlo rollouts using the RL policy network p_θ , but using 15,000 times less computation.

AlphaGO Zero (2017)

a Self-play



b Neural network training



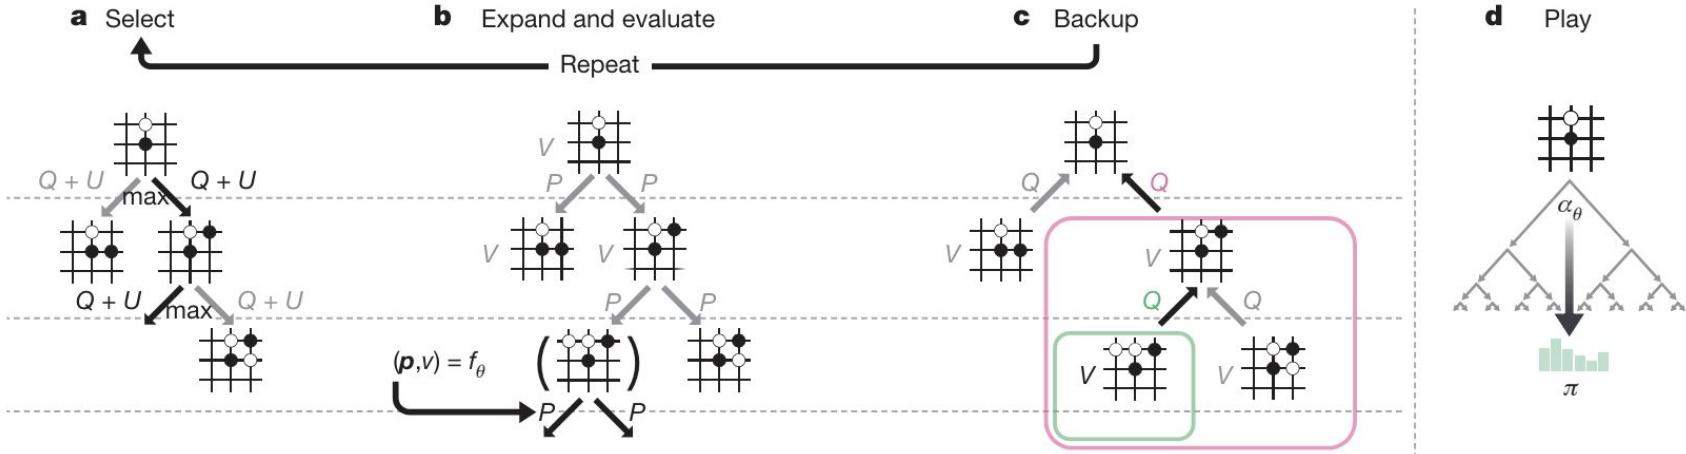


Figure 2 | MCTS in AlphaGo Zero. **a**, Each simulation traverses the tree by selecting the edge with maximum action value Q , plus an upper confidence bound U that depends on a stored prior probability P and visit count N for that edge (which is incremented once traversed). **b**, The leaf node is expanded and the associated position s is evaluated by the neural network $(P(s, \cdot), V(s)) = f_\theta(s)$; the vector of P values are stored in

the outgoing edges from s . **c**, Action value Q is updated to track the mean of all evaluations V in the subtree below that action. **d**, Once the search is complete, search probabilities π are returned, proportional to $N^{1/\tau}$, where N is the visit count of each move from the root state and τ is a parameter controlling temperature.

Self-play. The best current player α_{θ_*} , as selected by the evaluator, is used to generate data. In each iteration, α_{θ_*} plays 25,000 games of self-play, using 1,600 simulations of MCTS to select each move (this requires approximately 0.4 s per search). For the first 30 moves of each game, the temperature is set to $\tau = 1$; this selects moves proportionally to their visit count in MCTS, and ensures a diverse set of positions are encountered. For the remainder of the game, an infinitesimal temperature is used, $\tau \rightarrow 0$. Additional exploration is achieved by adding Dirichlet

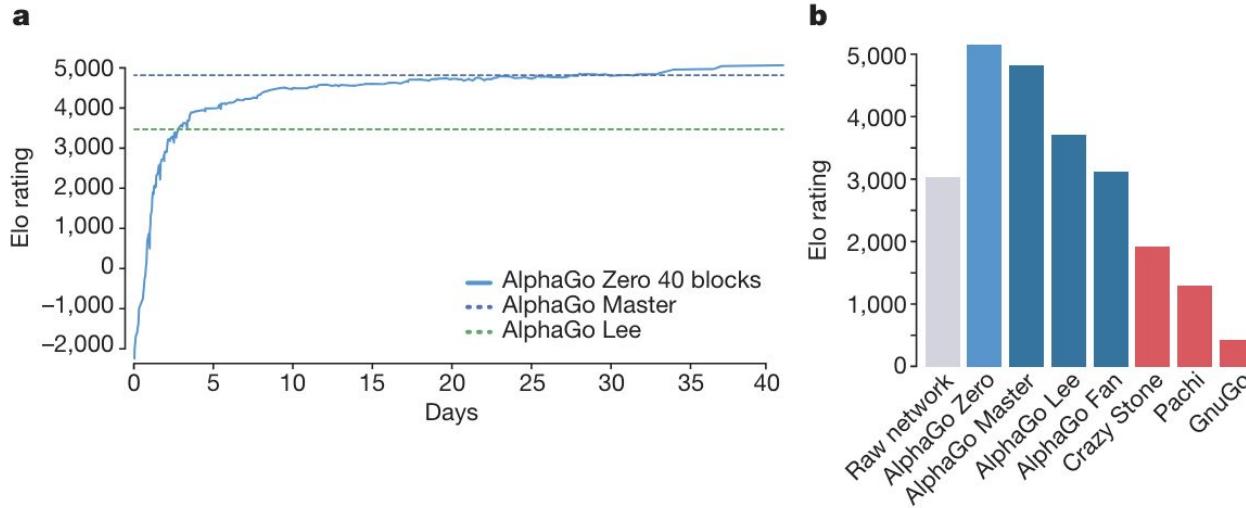


Figure 6 | Performance of AlphaGo Zero. **a**, Learning curve for AlphaGo Zero using a larger 40-block residual network over 40 days. The plot shows the performance of each player α_{θ_i} from each iteration i of our reinforcement learning algorithm. Elo ratings were computed from evaluation games between different players, using 0.4 s per search (see Methods). **b**, Final performance of AlphaGo Zero. AlphaGo Zero was trained for 40 days using a 40-block residual neural network. The plot shows the results of a tournament between: AlphaGo Zero, AlphaGo Master (defeated top human professionals 60–0 in online games), AlphaGo

Lee (defeated Lee Sedol), AlphaGo Fan (defeated Fan Hui), as well as previous Go programs Crazy Stone, Pachi and GnuGo. Each program was given 5 s of thinking time per move. AlphaGo Zero and AlphaGo Master played on a single machine on the Google Cloud; AlphaGo Fan and AlphaGo Lee were distributed over many machines. The raw neural network from AlphaGo Zero is also included, which directly selects the move a with maximum probability p_a , without using MCTS. Programs were evaluated on an Elo scale²⁵: a 200-point gap corresponds to a 75% probability of winning.



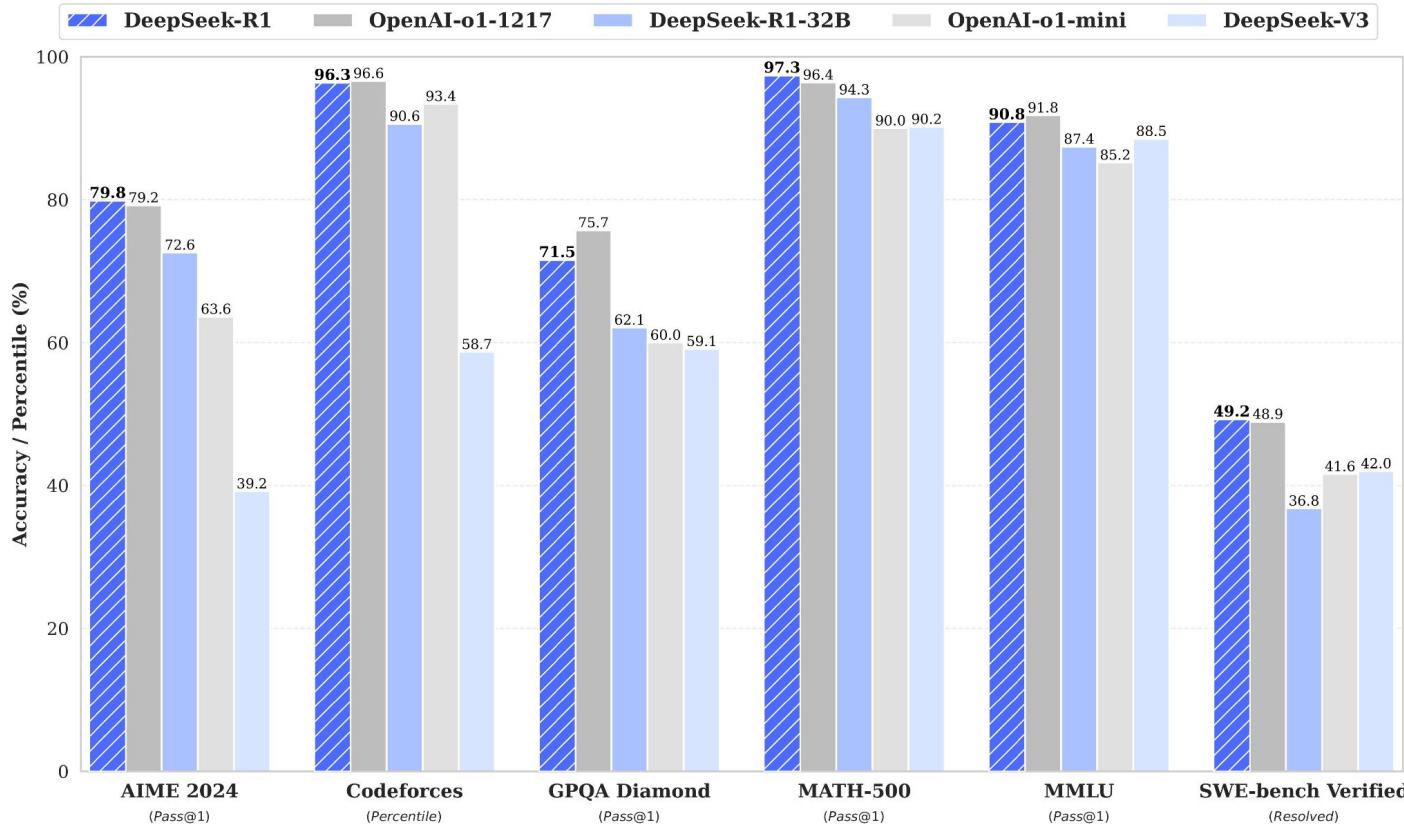
DeepSeek R1

Summary of DeepSeek-R1:

- **Overview:** Launched in January 2025 by Chinese AI startup DeepSeek, **DeepSeek-R1 is an open-source large language model** excelling in advanced reasoning tasks like math, coding, and logic, competing with OpenAI's o1.
- **Technology:** Built on the 671-billion-parameter DeepSeek-V3-Base, it uses **reinforcement learning (RL)** with minimal supervised fine-tuning (SFT) and **Group Relative Policy Optimization (GRPO)** for efficient training, achieving high performance (e.g., 79.8% on AIME, 97.3% on MATH-500).
- **Cost Efficiency:** Trained for ~\$6 million using ~2,000 Nvidia H800 chips, far less than the \$100 million–\$1 billion spent by U.S. competitors. API pricing is significantly lower (\$0.55/million input tokens vs. OpenAI's \$15).
- **Accessibility:** Available under MIT License, supporting commercial use and model distillation, with six smaller distilled models (1.5B–70B parameters). Powers DeepSeek's chatbot via web, app, and API.
- **Performance:** Outperforms or matches U.S. models like OpenAI's o1 and Meta's Llama in benchmarks, with a top-five ranking on Chatbot Arena.

Impact on Nvidia Stock (January–May 2025):

- **Initial Market Shock:** DeepSeek-R1's release on January 20, 2025, triggered a 17% drop in Nvidia's stock on January 27, erasing ~\$593 billion in market value, the largest single-day loss in Wall Street history. The cost-efficient model raised fears of reduced demand for Nvidia's high-end GPUs.



2.2.1. Reinforcement Learning Algorithm



Core technology of DeepSeek

Group Relative Policy Optimization In order to save the training costs of RL, we adopt Group Relative Policy Optimization (GRPO) (Shao et al., 2024), which foregoes the critic model that is typically the same size as the policy model, and estimates the baseline from group scores instead. Specifically, for each question q , GRPO samples a group of outputs $\{o_1, o_2, \dots, o_G\}$ from the old policy $\pi_{\theta_{old}}$ and then optimizes the policy model π_θ by maximizing the following objective:

$$\mathcal{J}_{GRPO}(\theta) = \mathbb{E}[q \sim P(Q), \{o_i\}_{i=1}^G \sim \pi_{\theta_{old}}(O|q)] \frac{1}{G} \sum_{i=1}^G \left(\min \left(\frac{\pi_\theta(o_i|q)}{\pi_{\theta_{old}}(o_i|q)} A_i, \text{clip} \left(\frac{\pi_\theta(o_i|q)}{\pi_{\theta_{old}}(o_i|q)}, 1 - \varepsilon, 1 + \varepsilon \right) A_i \right) - \beta \mathbb{D}_{KL}(\pi_\theta || \pi_{ref}) \right), \quad (1)$$

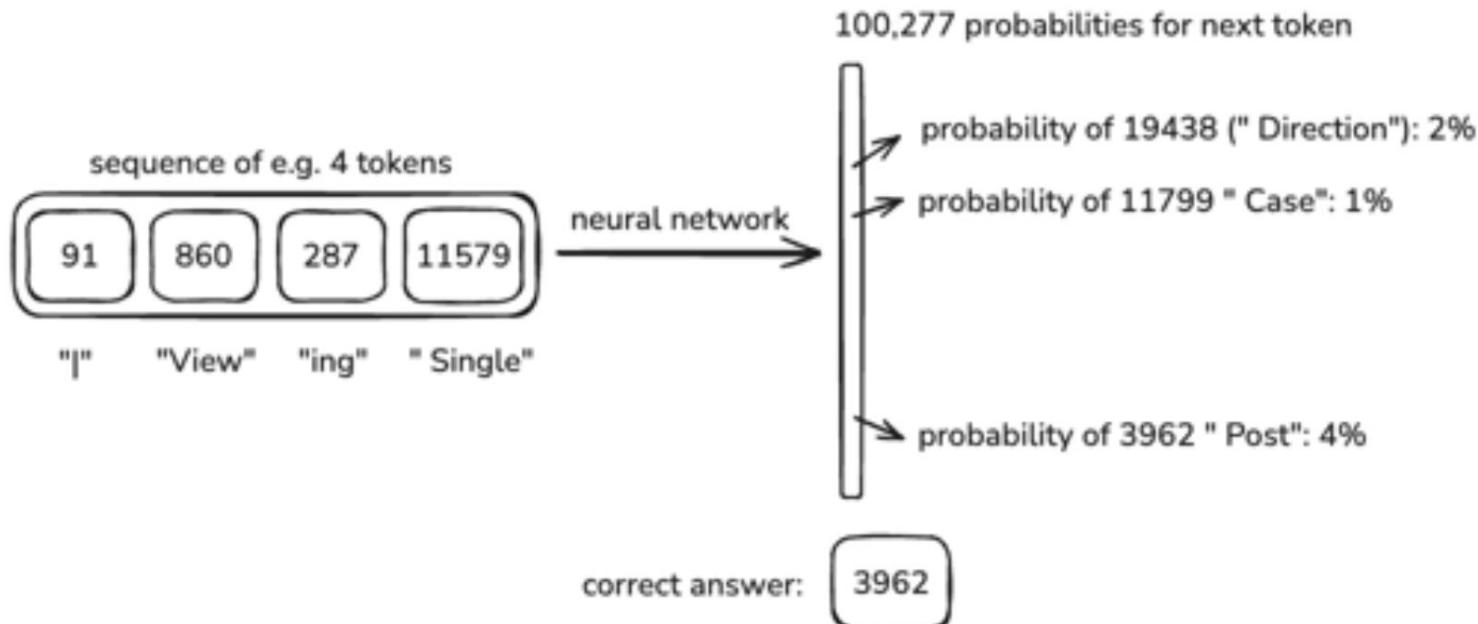
$$\mathbb{D}_{KL}(\pi_\theta || \pi_{ref}) = \frac{\pi_{ref}(o_i|q)}{\pi_\theta(o_i|q)} - \log \frac{\pi_{ref}(o_i|q)}{\pi_\theta(o_i|q)} - 1, \quad (2)$$

where ε and β are hyper-parameters, and A_i is the advantage, computed using a group of rewards $\{r_1, r_2, \dots, r_G\}$ corresponding to the outputs within each group:

$$A_i = \frac{r_i - \text{mean}(\{r_1, r_2, \dots, r_G\})}{\text{std}(\{r_1, r_2, \dots, r_G\})}. \quad (3)$$

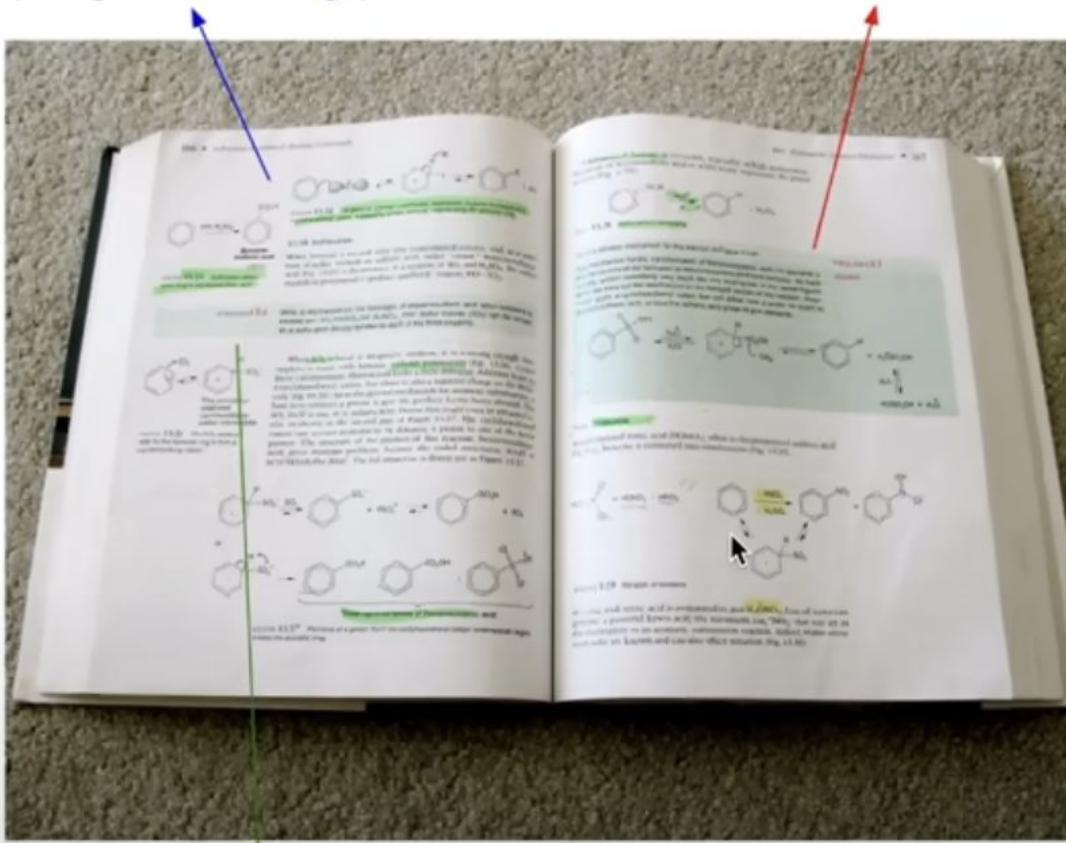
Next token generation as policy

Step 3: neural network training

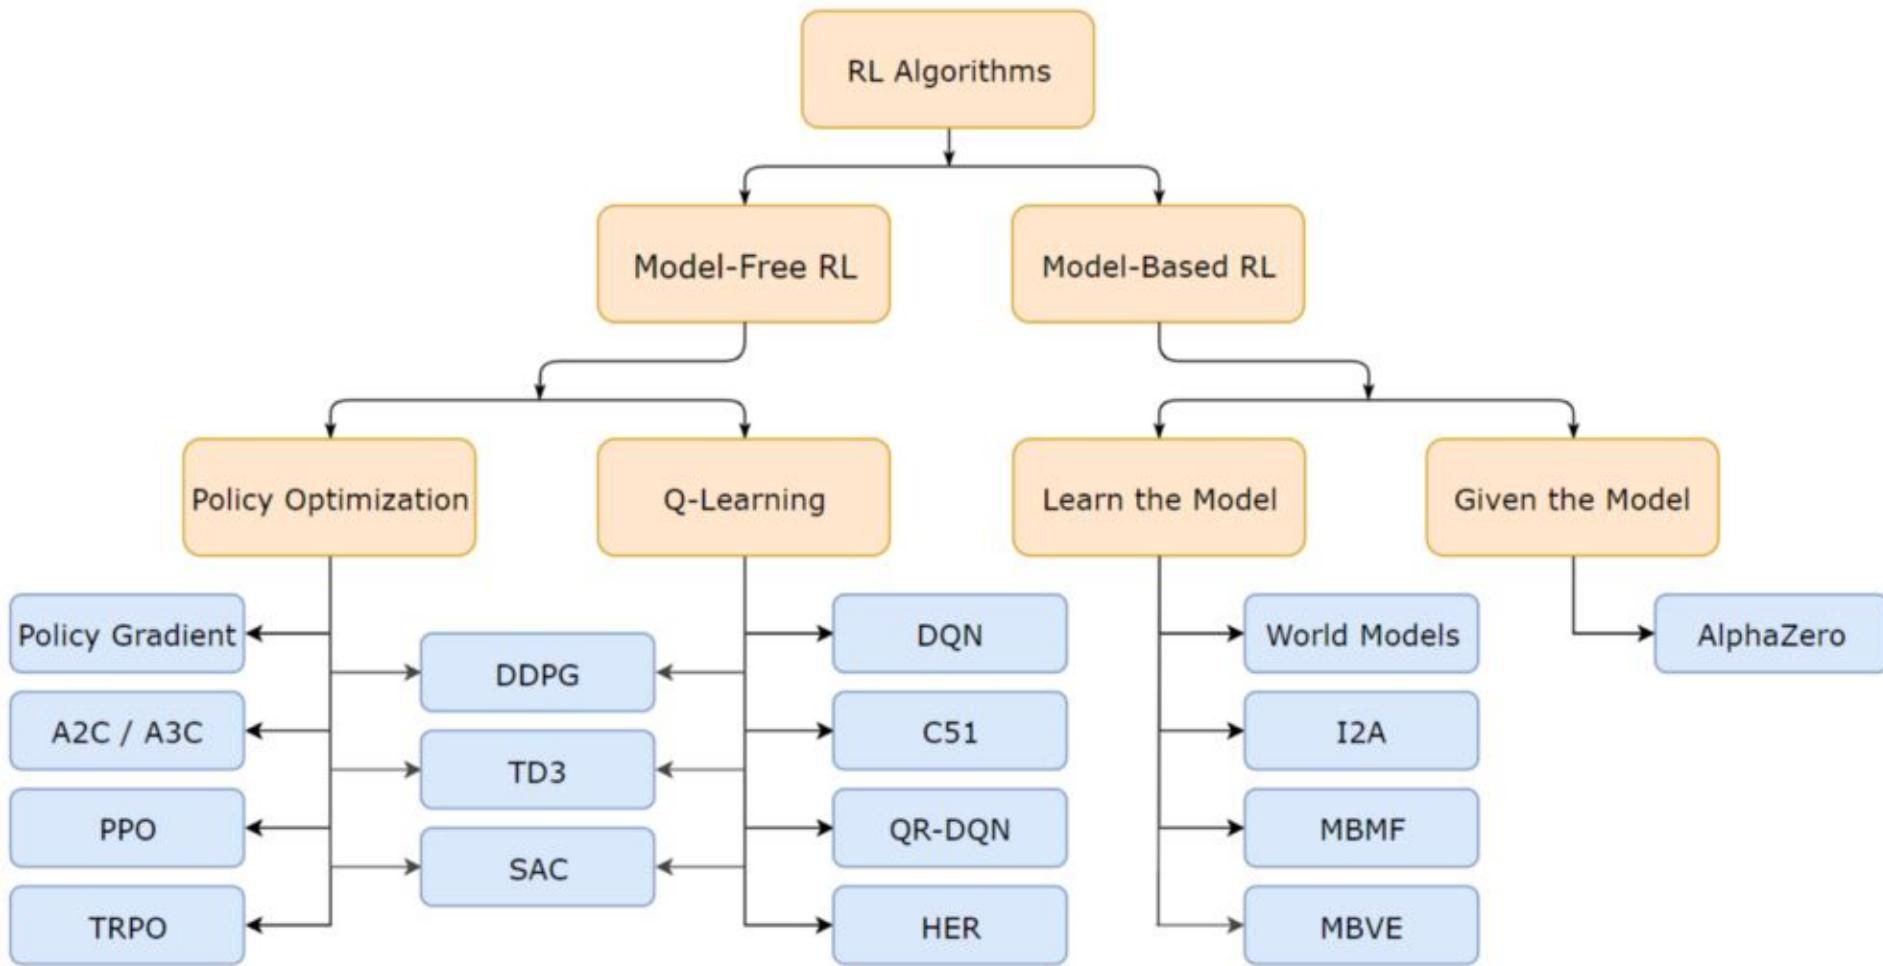


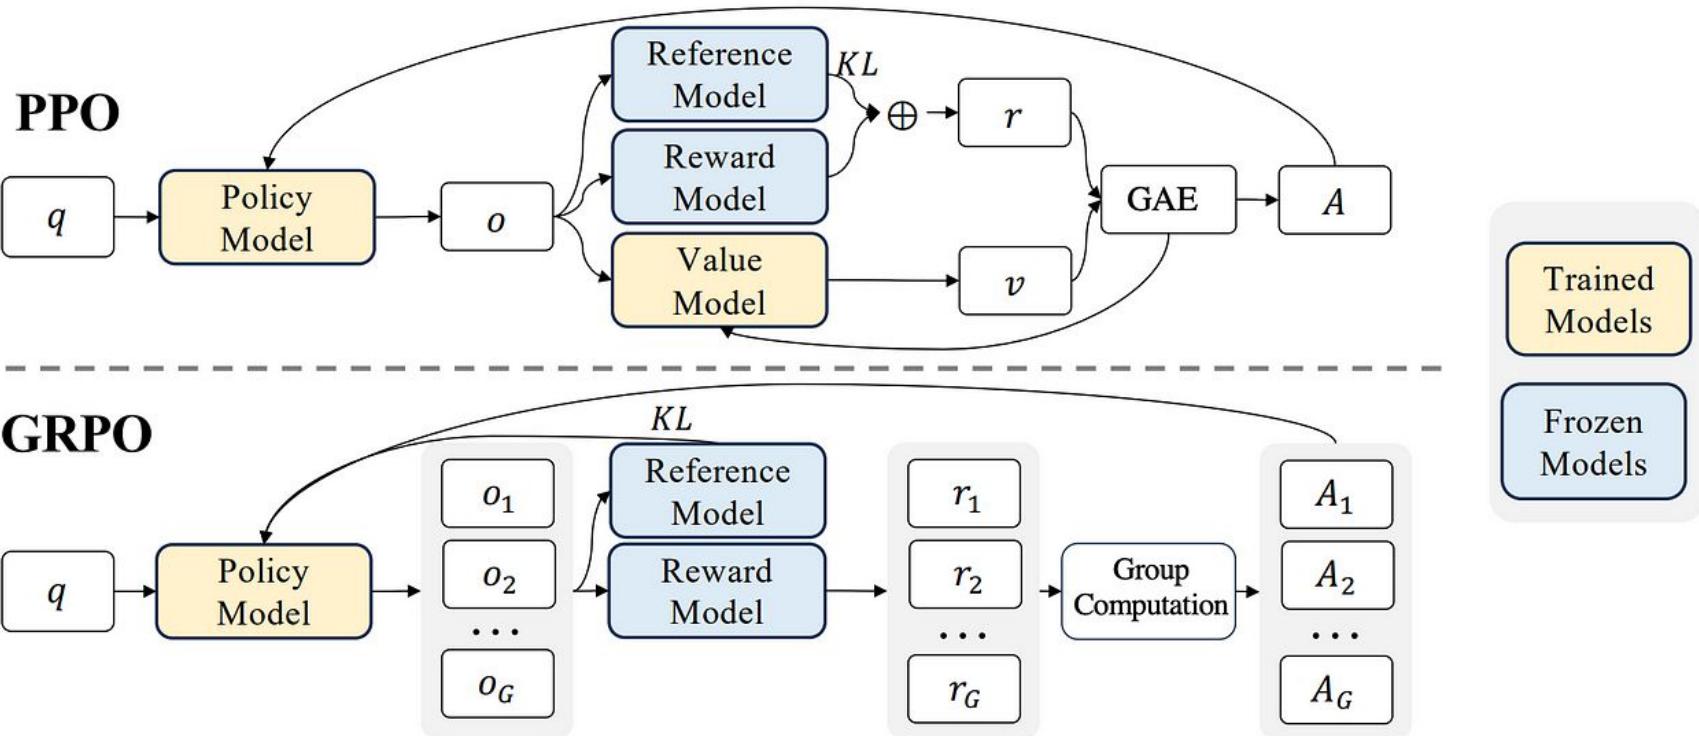
exposition \Leftrightarrow pretraining
(background knowledge)

worked problems \Leftrightarrow supervised finetuning
(problem + demonstrated solution, for imitation)



practice problems \Leftrightarrow reinforcement learning
(prompts to practice, trial & error until you reach the correct answer)





2.2.1. Reinforcement Learning Algorithm

Group Relative Policy Optimization In order to save the training costs of RL, we adopt Group Relative Policy Optimization (GRPO) (Shao et al., 2024), which foregoes the critic model that is typically the same size as the policy model, and estimates the baseline from group scores instead. Specifically, for each question q , GRPO samples a group of outputs $\{o_1, o_2, \dots, o_G\}$ from the old policy $\pi_{\theta_{old}}$ and then optimizes the policy model π_θ by maximizing the following objective:

$$\mathcal{J}_{GRPO}(\theta) = \mathbb{E}[q \sim P(Q), \{o_i\}_{i=1}^G \sim \pi_{\theta_{old}}(O|q)] \left(\frac{1}{G} \sum_{i=1}^G \left(\min \left(\frac{\pi_\theta(o_i|q)}{\pi_{\theta_{old}}(o_i|q)} A_i, \text{clip} \left(\frac{\pi_\theta(o_i|q)}{\pi_{\theta_{old}}(o_i|q)}, 1 - \varepsilon, 1 + \varepsilon \right) A_i \right) - \beta \mathbb{D}_{KL}(\pi_\theta || \pi_{ref}) \right) \right), \quad (1)$$

$$\mathbb{D}_{KL}(\pi_\theta || \pi_{ref}) = \frac{\pi_{ref}(o_i|q)}{\pi_\theta(o_i|q)} - \log \frac{\pi_{ref}(o_i|q)}{\pi_\theta(o_i|q)} - 1, \quad (2)$$

where ε and β are hyper-parameters, and A_i is the advantage, computed using a group of rewards $\{r_1, r_2, \dots, r_G\}$ corresponding to the outputs within each group:

$$A_i = \frac{r_i - \text{mean}(\{r_1, r_2, \dots, r_G\})}{\text{std}(\{r_1, r_2, \dots, r_G\})}. \quad (3)$$

Proximal Policy Optimization Algorithms

John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, Oleg Klimov

OpenAI

{joschu, filip, prafulla, alec, oleg}@openai.com

Abstract

We propose a new family of policy gradient methods for reinforcement learning, which alternate between sampling data through interaction with the environment, and optimizing a “surrogate” objective function using stochastic gradient ascent. Whereas standard policy gradient methods perform one gradient update per data sample, we propose a novel objective function that enables multiple epochs of minibatch updates. The new methods, which we call proximal policy optimization (PPO), have some of the benefits of trust region policy optimization (TRPO), but they are much simpler to implement, more general, and have better sample complexity (empirically). Our experiments test PPO on a collection of benchmark tasks, including simulated robotic locomotion and Atari game playing, and we show that PPO outperforms other online policy gradient methods, and overall strikes a favorable balance between sample complexity, simplicity, and wall-time.

2.1 Policy Gradient Methods

Policy gradient methods work by computing an estimator of the policy gradient and plugging it into a stochastic gradient ascent algorithm. The most commonly used gradient estimator has the form

$$\hat{g} = \hat{\mathbb{E}}_t \left[\nabla_{\theta} \log \pi_{\theta}(a_t | s_t) \hat{A}_t \right] \quad (1)$$

where π_{θ} is a stochastic policy and \hat{A}_t is an estimator of the advantage function at timestep t . Here, the expectation $\hat{\mathbb{E}}_t[\dots]$ indicates the empirical average over a finite batch of samples, in an algorithm that alternates between sampling and optimization. Implementations that use automatic differentiation software work by constructing an objective function whose gradient is the policy gradient estimator; the estimator \hat{g} is obtained by differentiating the objective

$$L^{PG}(\theta) = \hat{\mathbb{E}}_t \left[\log \pi_{\theta}(a_t | s_t) \hat{A}_t \right]. \quad (2)$$

While it is appealing to perform multiple steps of optimization on this loss L^{PG} using the same trajectory, doing so is not well-justified, and empirically it often leads to destructively large policy updates (see Section 6.1; results are not shown but were similar or worse than the “no clipping or penalty” setting).

2.2 Trust Region Methods

In TRPO [Sch+15b], an objective function (the “surrogate” objective) is maximized subject to a constraint on the size of the policy update. Specifically,

$$\underset{\theta}{\text{maximize}} \quad \hat{\mathbb{E}}_t \left[\frac{\pi_{\theta}(a_t | s_t)}{\pi_{\theta_{\text{old}}}(a_t | s_t)} \hat{A}_t \right] \quad (3)$$

$$\text{subject to} \quad \hat{\mathbb{E}}_t [\text{KL}[\pi_{\theta_{\text{old}}}(\cdot | s_t), \pi_{\theta}(\cdot | s_t)]] \leq \delta. \quad (4)$$

Here, θ_{old} is the vector of policy parameters before the update. This problem can efficiently be approximately solved using the conjugate gradient algorithm, after making a linear approximation to the objective and a quadratic approximation to the constraint.

The theory justifying TRPO actually suggests using a penalty instead of a constraint, i.e., solving the unconstrained optimization problem

$$\underset{\theta}{\text{maximize}} \hat{\mathbb{E}}_t \left[\frac{\pi_{\theta}(a_t | s_t)}{\pi_{\theta_{\text{old}}}(a_t | s_t)} \hat{A}_t - \beta \text{KL}[\pi_{\theta_{\text{old}}}(\cdot | s_t), \pi_{\theta}(\cdot | s_t)] \right] \quad (5)$$

3 Clipped Surrogate Objective

Let $r_t(\theta)$ denote the probability ratio $r_t(\theta) = \frac{\pi_\theta(a_t | s_t)}{\pi_{\theta_{\text{old}}}(a_t | s_t)}$, so $r(\theta_{\text{old}}) = 1$. TRPO maximizes a “surrogate” objective

$$L^{CPI}(\theta) = \hat{\mathbb{E}}_t \left[\frac{\pi_\theta(a_t | s_t)}{\pi_{\theta_{\text{old}}}(a_t | s_t)} \hat{A}_t \right] = \hat{\mathbb{E}}_t \left[r_t(\theta) \hat{A}_t \right]. \quad (6)$$

The superscript *CPI* refers to conservative policy iteration [KL02], where this objective was proposed. Without a constraint, maximization of L^{CPI} would lead to an excessively large policy update; hence, we now consider how to modify the objective, to penalize changes to the policy that move $r_t(\theta)$ away from 1.

The main objective we propose is the following:

$$L^{CLIP}(\theta) = \hat{\mathbb{E}}_t \left[\min(r_t(\theta) \hat{A}_t, \text{clip}(r_t(\theta), 1 - \epsilon, 1 + \epsilon) \hat{A}_t) \right] \quad (7)$$

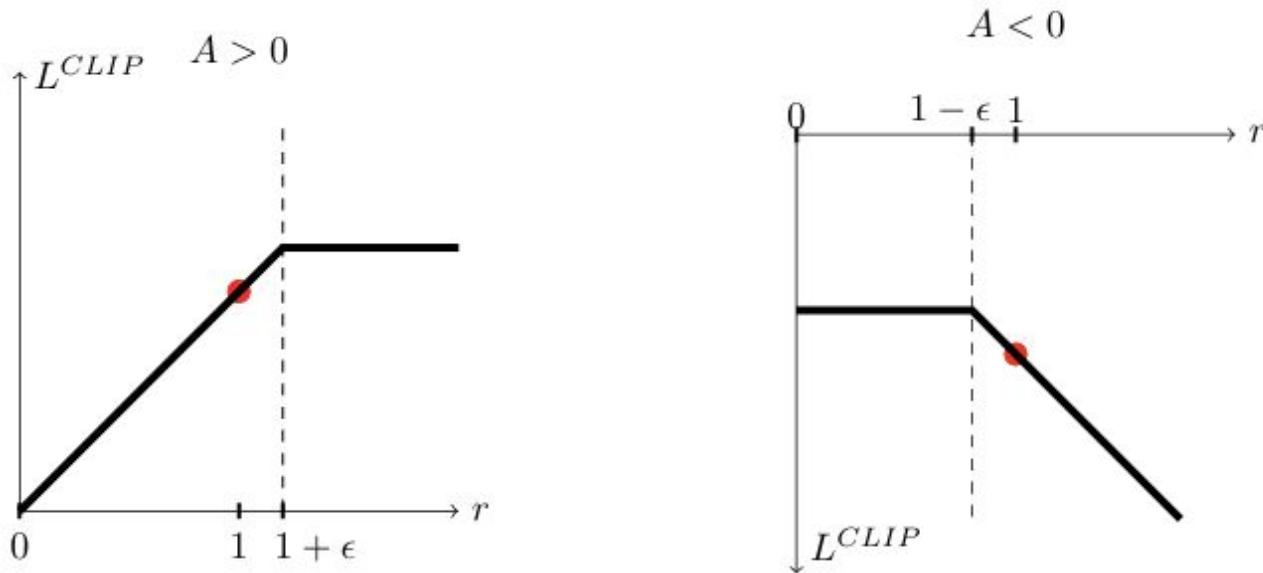


Figure 1: Plots showing one term (i.e., a single timestep) of the surrogate function L^{CLIP} as a function of the probability ratio r , for positive advantages (left) and negative advantages (right). The red circle on each plot shows the starting point for the optimization, i.e., $r = 1$. Note that L^{CLIP} sums many of these terms.

DeepResearch

Good evening, Joshua Yao-Yu.
How can I help you today?

What do you want to know?



DeepSearch



Think

Grok 3



Hello, Joshua Yao-Yu

Tell me what
you can do

Save me
time

Research
a topic

Write an essay
on the history of chess

Ask Gemini



Deep Research



Canvas



AlphaEvolve



- Advancing the frontiers in mathematics and algorithm discovery
- Enhancing AI training and inference
 - AlphaEvolve achieved up to a 32.5% speedup for the [FlashAttention](#) kernel implementation in [Transformer](#)-based AI models
- Designing better algorithms with large language models



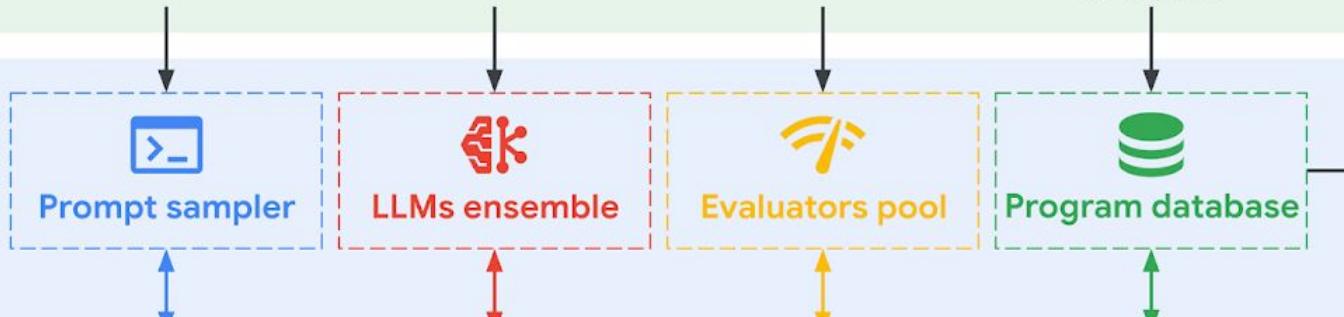
Scientist / Engineer

Prompt template and configuration

Choice of existing or custom LLMs

Evaluation code

Initial program with components to evolve



Best program

Distributed Controller Loop

```
parent_program, inspirations = database.sample()  
prompt = prompt_sampler.build(parent_program, inspirations)  
diff = llm.generate(prompt)  
child_program = apply_diff(parent_program, diff)  
results = evaluator.execute(child_program)  
database.add(child_program, results)
```



AlphaEvolve

$$\max_{-1/2 \leq t \leq 1/2} \int_{\mathbb{R}} f(t-x) f(x) dx \geq \mathbb{C} \left(\int_{-1/4}^{1/4} f(x) dx \right)^2$$

$$1.5098 \rightarrow \mathbf{1.5053}$$

$$\|f * f\|_2^2 \leq \mathbb{C}' \|f * f\|_1 \|f * f\|_\infty$$

$$0.8892 \rightarrow \mathbf{0.8962}$$

$$\max_{-1/2 \leq t \leq 1/2} \left| \int_{\mathbb{R}} f(t-x) f(x) dx \right| \geq \mathbb{C}'' \left(\int_{-1/4}^{1/4} f(x) dx \right)^2$$

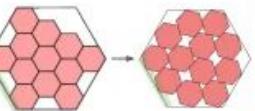
$$1.4581 \rightarrow \mathbf{1.4557}$$

$$A(f)A(\hat{f}) \geq \mathbb{C}'''$$

$$0.3523 \rightarrow \mathbf{0.3521}$$

Analysis

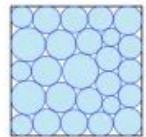
Hexagon outer edge
4.000 \rightarrow **3.942**



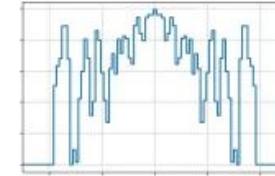
Max distance/min distance
12.890 \rightarrow **12.889**



Sum of radii
2.6340 \rightarrow **2.6358**



Geometry



$$\sup_{x \in [-2, 2]} \int_{-1}^1 f(t)g(x+t) dt \geq \mathbb{C}$$

$$0.380926 \rightarrow \mathbf{0.380924}$$

$$|A+B| \ll |A|$$

$$|A-B| \gg |A|^{\mathbb{C}}$$

$$1.1446 \rightarrow \mathbf{1.1584}$$

Combinatorics

The Strassen algorithm partitions A , B and C into equally sized **block matrices**

$$A = \begin{bmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{bmatrix}, \quad B = \begin{bmatrix} B_{11} & B_{12} \\ B_{21} & B_{22} \end{bmatrix}, \quad C = \begin{bmatrix} C_{11} & C_{12} \\ C_{21} & C_{22} \end{bmatrix},$$

with $A_{ij}, B_{ij}, C_{ij} \in \text{Mat}_{2^{n-1} \times 2^{n-1}}(\mathcal{R})$. The naive algorithm would be:

$$\begin{bmatrix} C_{11} & C_{12} \\ C_{21} & C_{22} \end{bmatrix} = \begin{bmatrix} A_{11} \times B_{11} + A_{12} \times B_{21} & A_{11} \times B_{12} + A_{12} \times B_{22} \\ A_{21} \times B_{11} + A_{22} \times B_{21} & A_{21} \times B_{12} + A_{22} \times B_{22} \end{bmatrix}.$$

This construction does not reduce the number of multiplications: 8 multiplications of matrix blocks are still needed to calculate the C_{ij} matrices, the same number of multiplications needed when using standard matrix multiplication.

The Strassen algorithm defines instead new values:

$$M_1 = (A_{11} + A_{22}) \times (B_{11} + B_{22});$$

$$M_2 = (A_{21} + A_{22}) \times B_{11};$$

$$M_3 = A_{11} \times (B_{12} - B_{22});$$

$$M_4 = A_{22} \times (B_{21} - B_{11});$$

$$M_5 = (A_{11} + A_{12}) \times B_{22};$$

$$M_6 = (A_{21} - A_{11}) \times (B_{11} + B_{12});$$

$$M_7 = (A_{12} - A_{22}) \times (B_{21} + B_{22}),$$

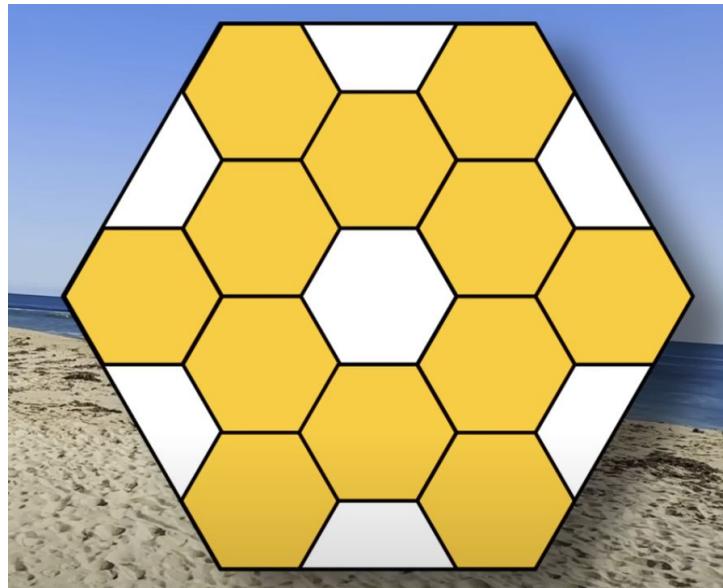
using only 7 multiplications (one for each M_k) instead of 8. We may now express the C_{ij} in terms of M_k :

$$\begin{bmatrix} C_{11} & C_{12} \\ C_{21} & C_{22} \end{bmatrix} = \begin{bmatrix} M_1 + M_4 - M_5 + M_7 & M_3 + M_5 \\ M_2 + M_4 & M_1 - M_2 + M_3 + M_6 \end{bmatrix}.$$

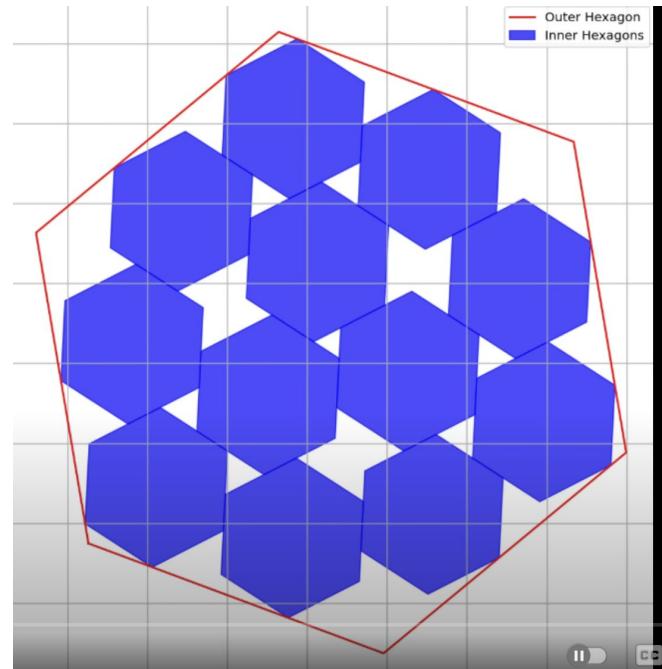
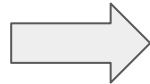
[Strassen's 1969 algorithm](#)

Packing problem: Hexagons in Hexagons

The following pictures show n regular hexagons with side 1 packed inside the smallest known regular hexagon (of side length s).



side lengths = 4



side lengths = 3.942

New result distribution

Visualization of results across 67 problems.

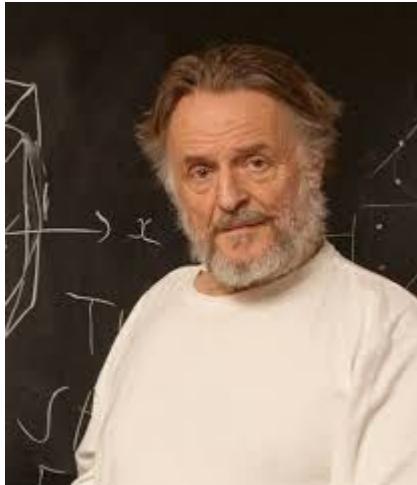
■ New result ■ Former new result, got improved upon ■ Worse than literature bound
■ Matched known optimal bound ■ Matched literature bound / N/A



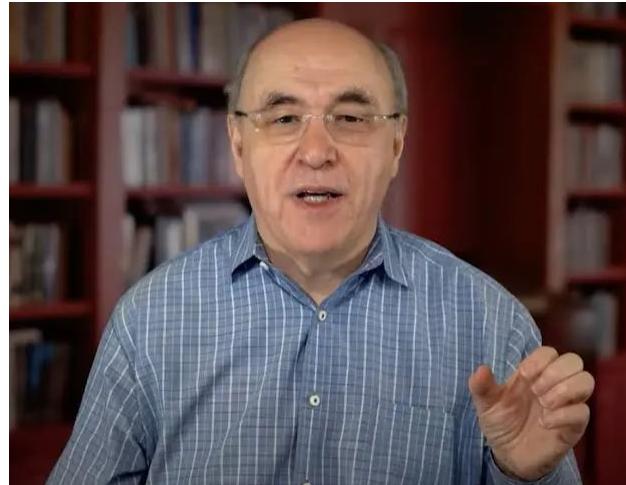
Terry Tao
Blogpost:
[Mathematical exploration
and discovery at scale](#)

Bonus (emergent Phenome):

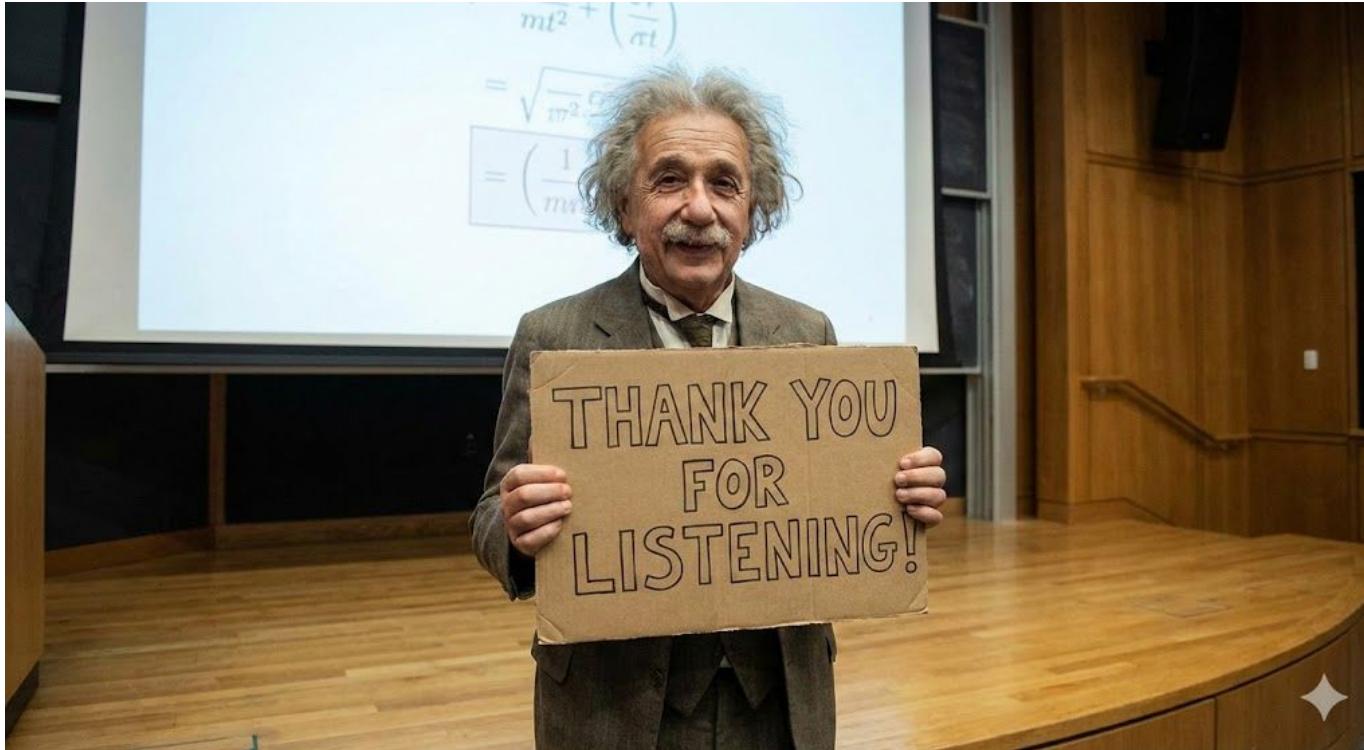
- [Game of Life](#)
- [Neural Cellular Automata](#)



John Conway



Stephen Wolfram



p.s. We're hiring! Please feel free to reach out!

Generated by Gemini3/Nanobanaba pro