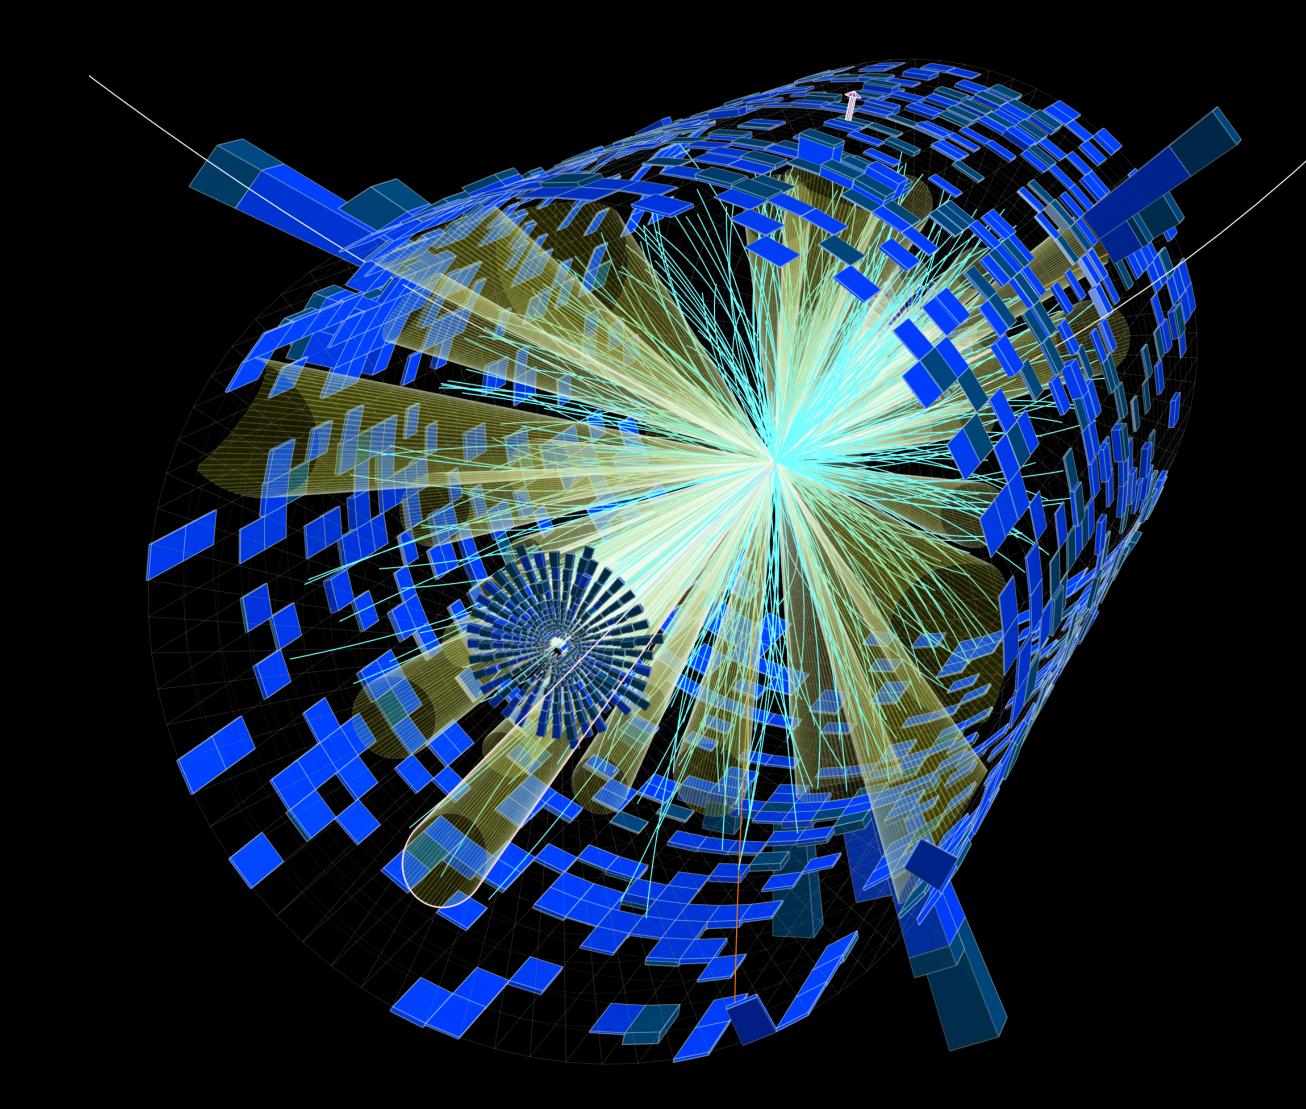
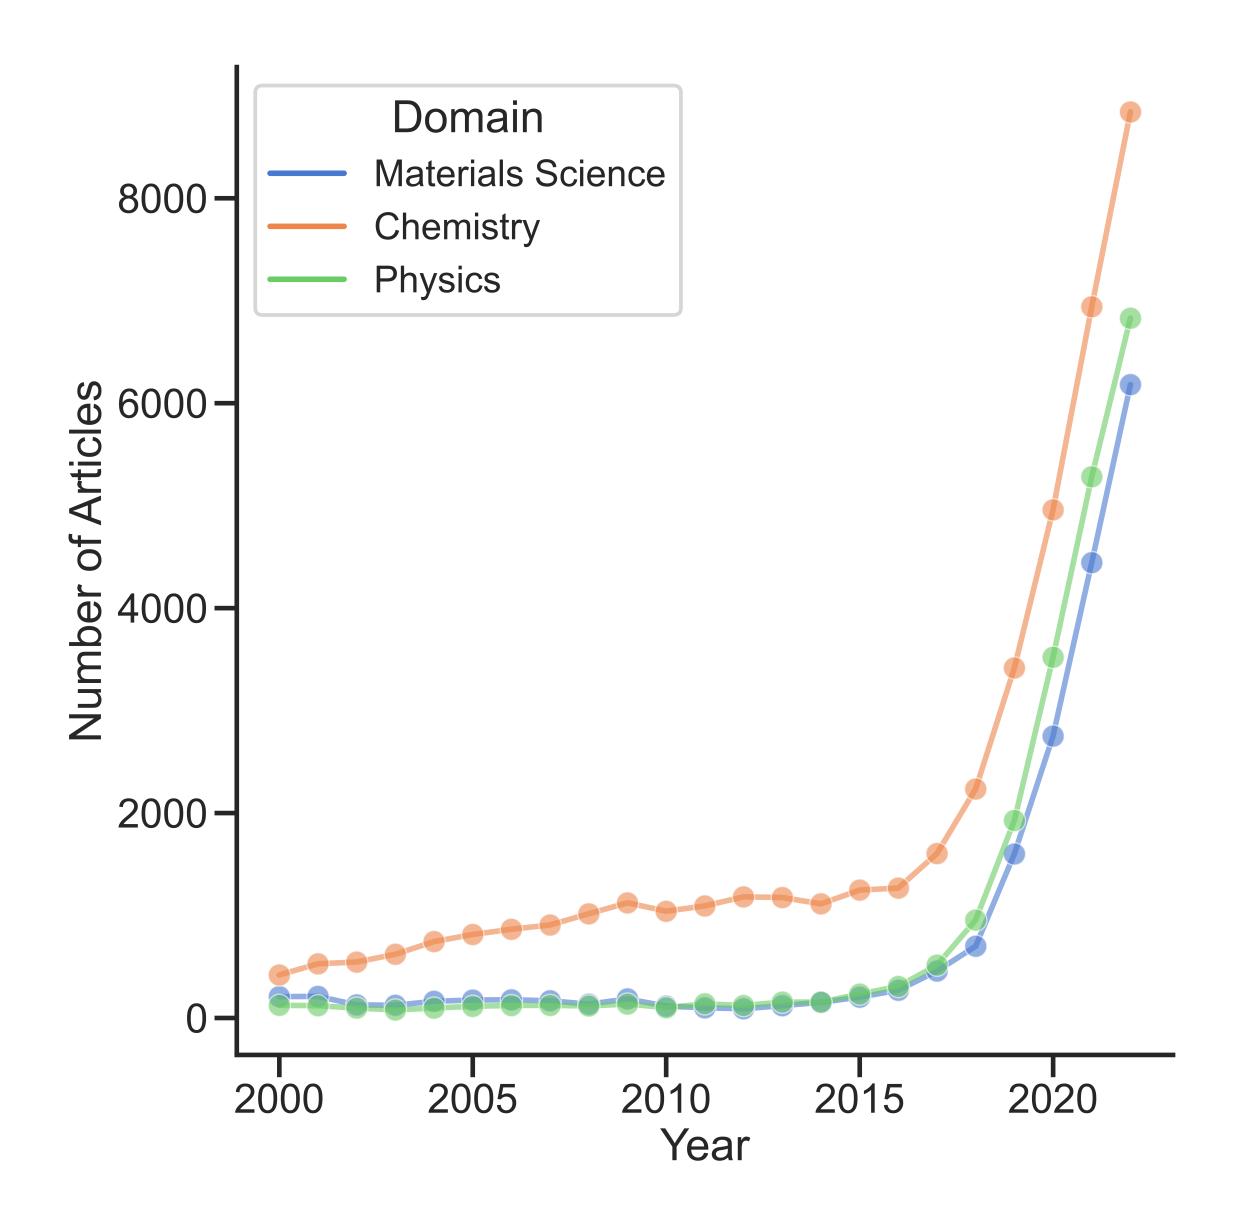
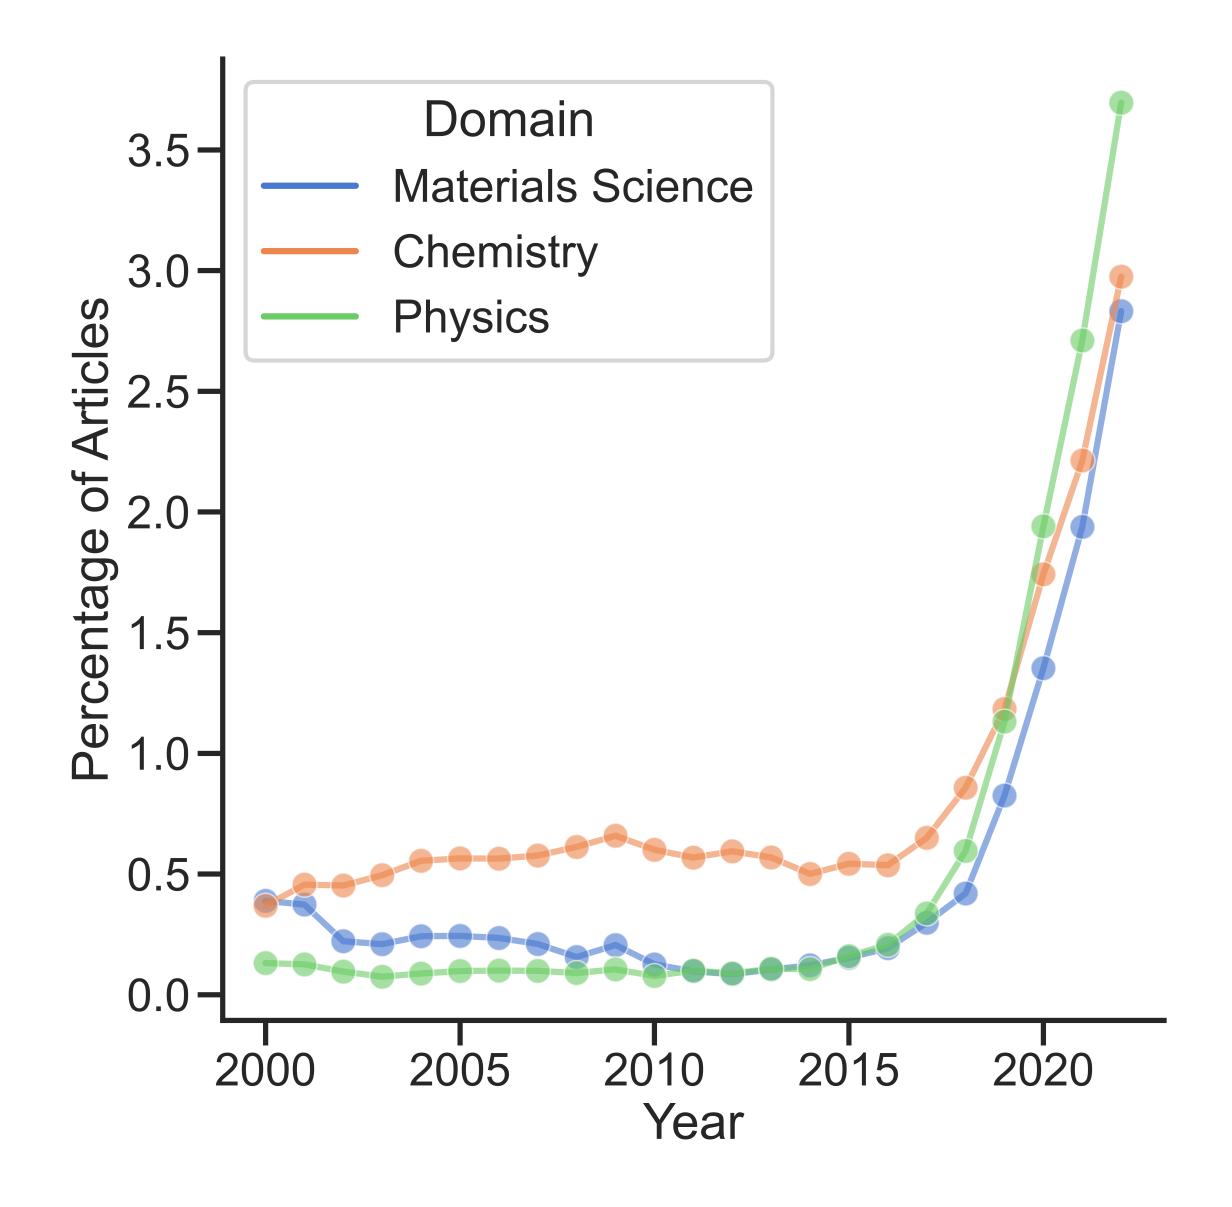
Al for Fundamental Physics





ML Publications in Science



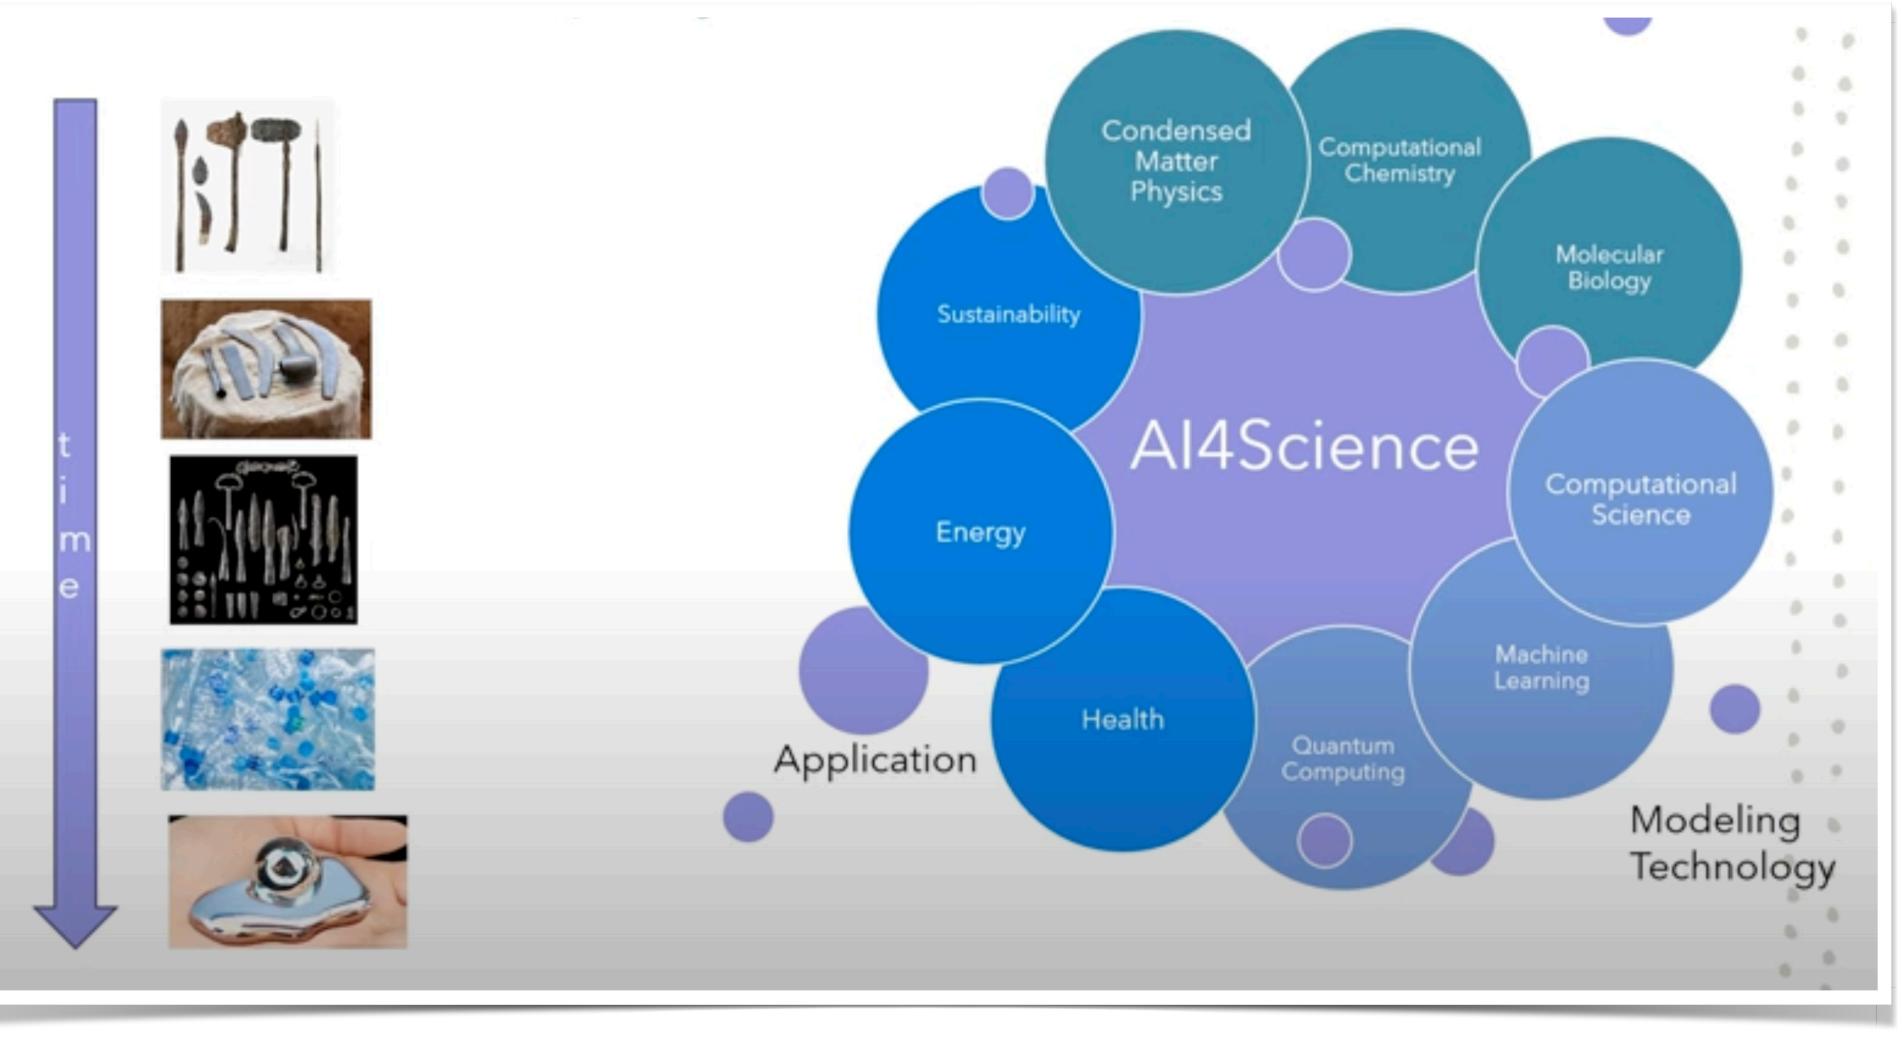


Max Welling's talk

"Where is fundamental physics in Al4Science?" — Lukas Heinrich

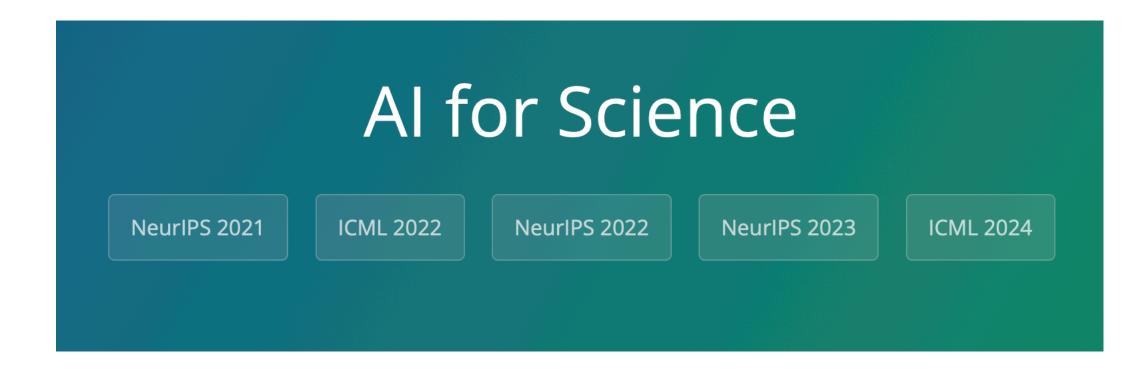
• Hint: was not referring to condensed matter physics



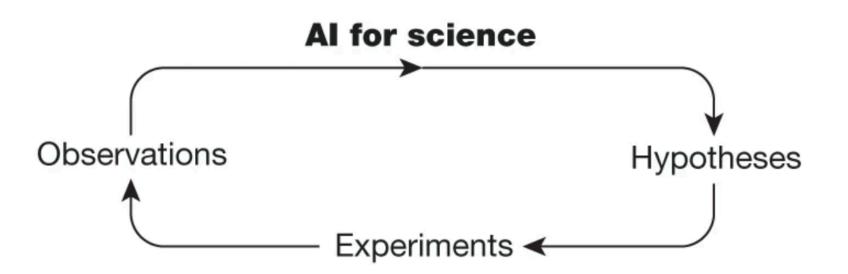


Al for Science Workshop Series

https://ai4sciencecommunity.github.io/









Weather forecasting



Rare event selection in particle collisions



Battery design optimization



Language modelling for biomedical sequences



Magnetic control of nuclear fusion reactors



High-throughput virtual screening



Planning chemical synthesis pathway



Navigation in the hypothesis space



Neural solvers of differential equations

Hydropower station

location planning



Super-resolution 3D live-cell imaging



Synthetic electronic health record generation



Symbolic regression

Al for Science Workshop Series

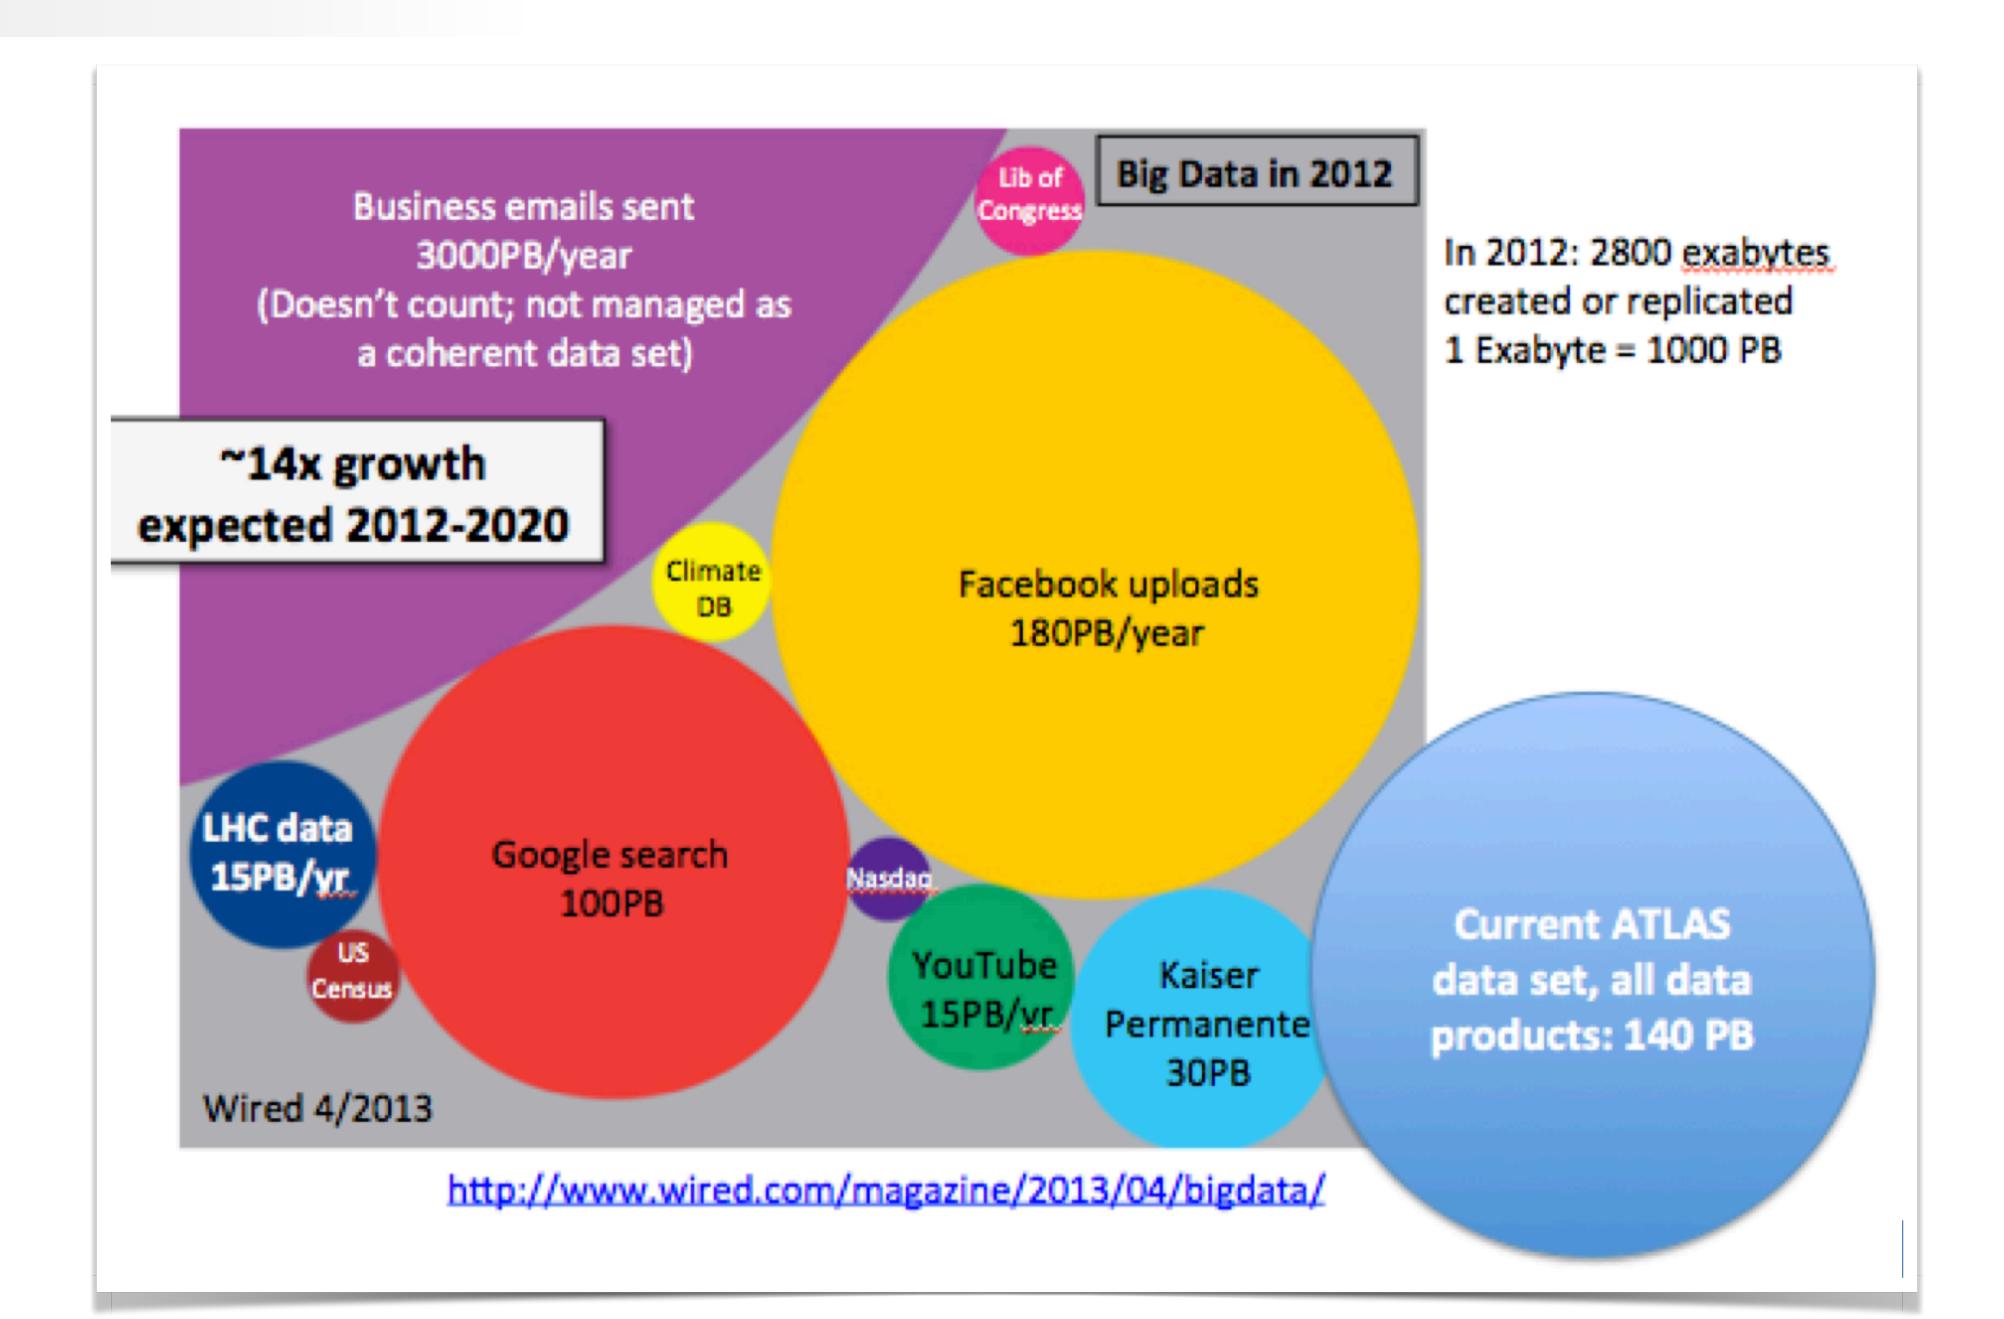
https://ai4sciencecommunity.github.io/





Invited Talk	Weinan E	8:15-9:10am
Invited Talk	Shuiwang Ji	9:10-10:05am
Coffee Break	-	10:05-10:15am
Invited Talk	Maria Schuld	10:15-11:10am
Invited Talk	David Baker	11:10-12:05pm
Poster Session		12:05pm-13:00pm
Contributed Talk	DiffDock: Diffusion Steps, Twists, and Turns for Molecular Docking	13:00-13:10pm
Contributed Talk	An "interpretable-by-design" neural network to decipher RNA splicing regulatory logic	13:10-13:20pm
Contributed Talk	FALCON: Fourier Adaptive Learning and Control for Disturbance Rejection Under Extreme Turbulence	13:20-13:30pm
Contributed Talk	Structural Causal Model for Molecular Dynamics Simulation	13:30-13:40pm
Contributed Talk	Structure-Inducing Pre-training	13:40-13:50pm
Panel Discussion	Panelists: Peter Lee, Sara-Jane Dunn, Arvind Rajpal, Alexander Rives, Matt Welborn Moderator: Max Welling	13:55-14:55pm
Coffee Break	-	14:55-15:10pm
Invited Talk	Jimeng Sun	15:10-16:05pm
Invited Talk	Tess Smidt	16:05-17:00pm

Compare with Big Data

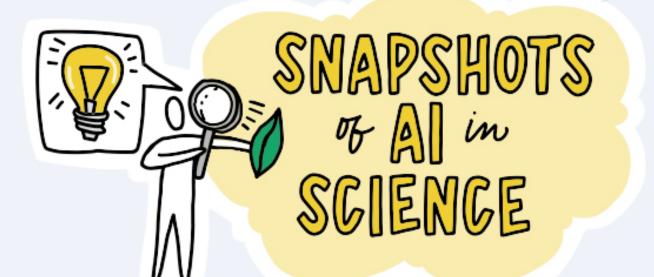


What are some of the dominant narratives about AI for Science?

How does fundamental physics fit in?

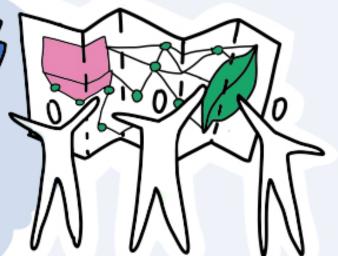
MACHINE LEARNING in SCIENCE

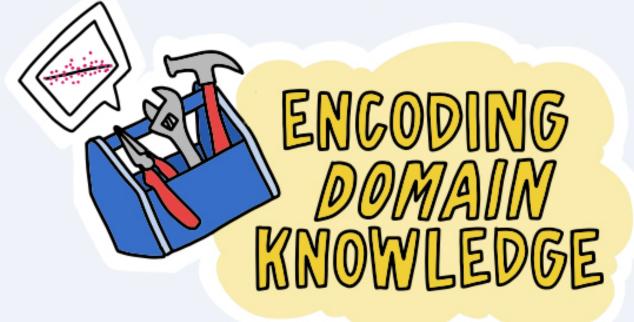
BRIDGING DATA-DRIVEN& MECHANISTIC MODELLING APPROACHES



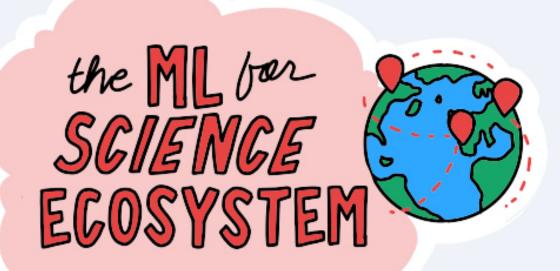


CONNECTING DATA to CAUSALITY

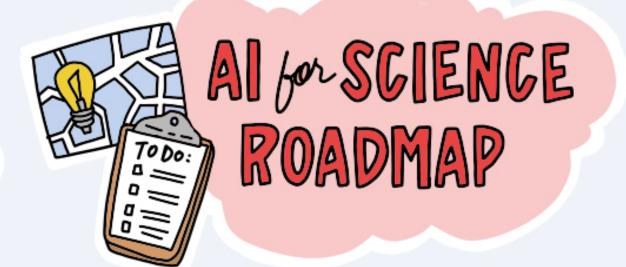












PARTICIPANTS'TALKS

LEARNING from RECENT EXPERIENCES

> CONNECTING DATA to CAUSALITY

BUILDING EFFECTIVE SIMULATIONS

ENCODING DOMAIN
KNOWLEDGE



CLICK EACH TITLE / to go to PAGE













FOURTH
PARADIGM

DATA-INTENSIVE SCIENTIFIC DISCOVERY

Scientific Paradigms:

- 1. Empirical
- 2. Theoretical
- 3. Computational
- 4. Data-Driven / Data-Intensive Discovery

5. AI/ML + Simulation + Data



Published July 7, 2022

By Christopher Bishop, Technical Fellow and Director, Microsoft Research Al4Science

hare this page

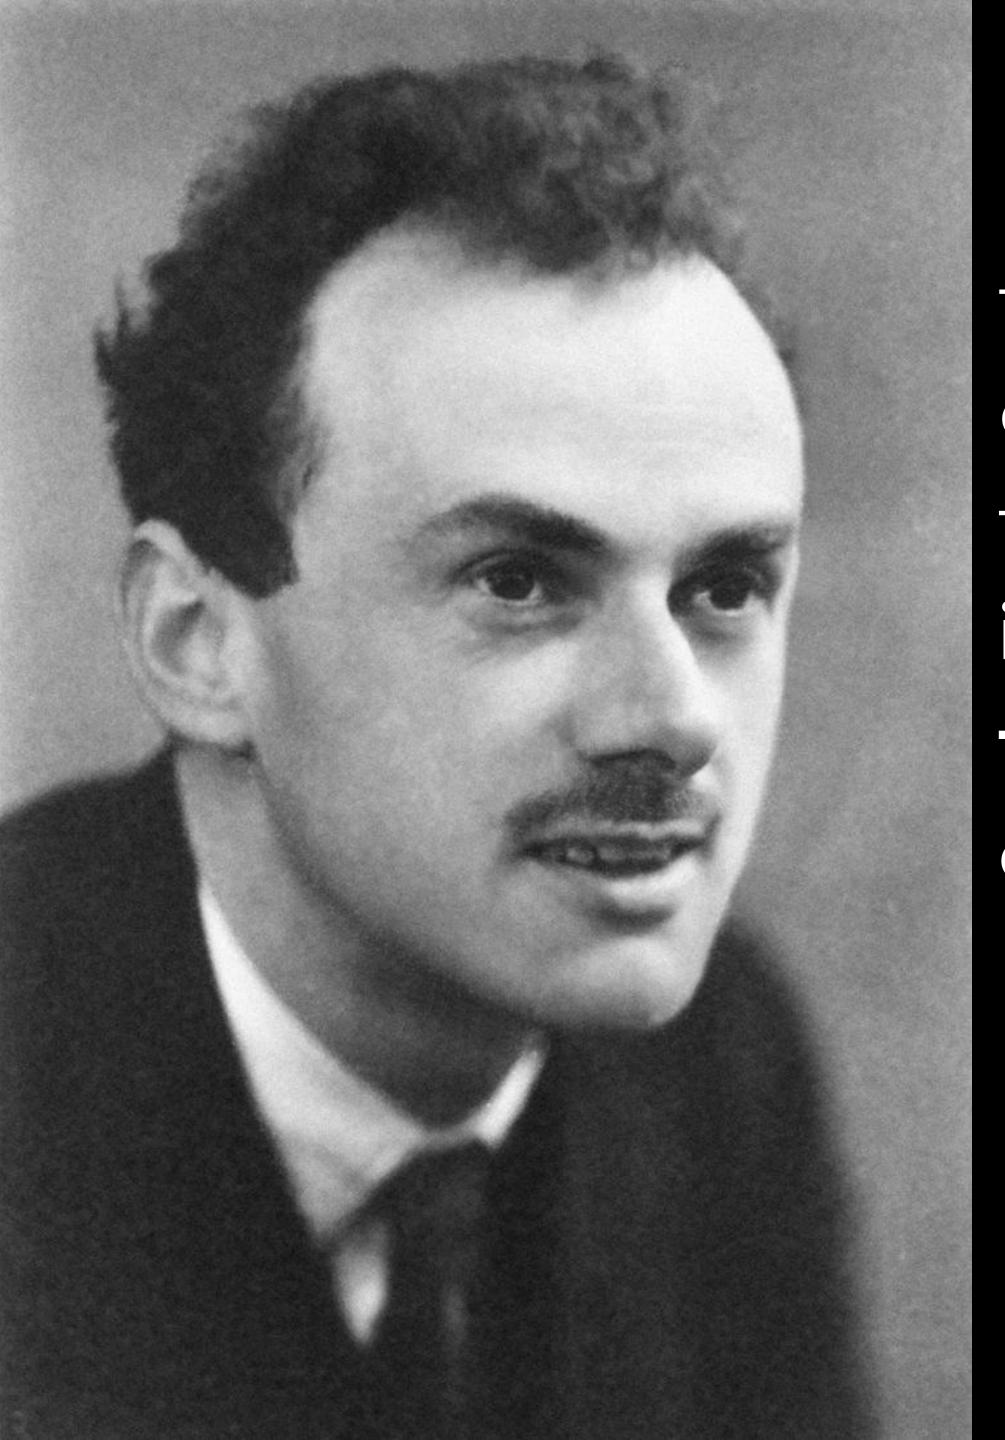












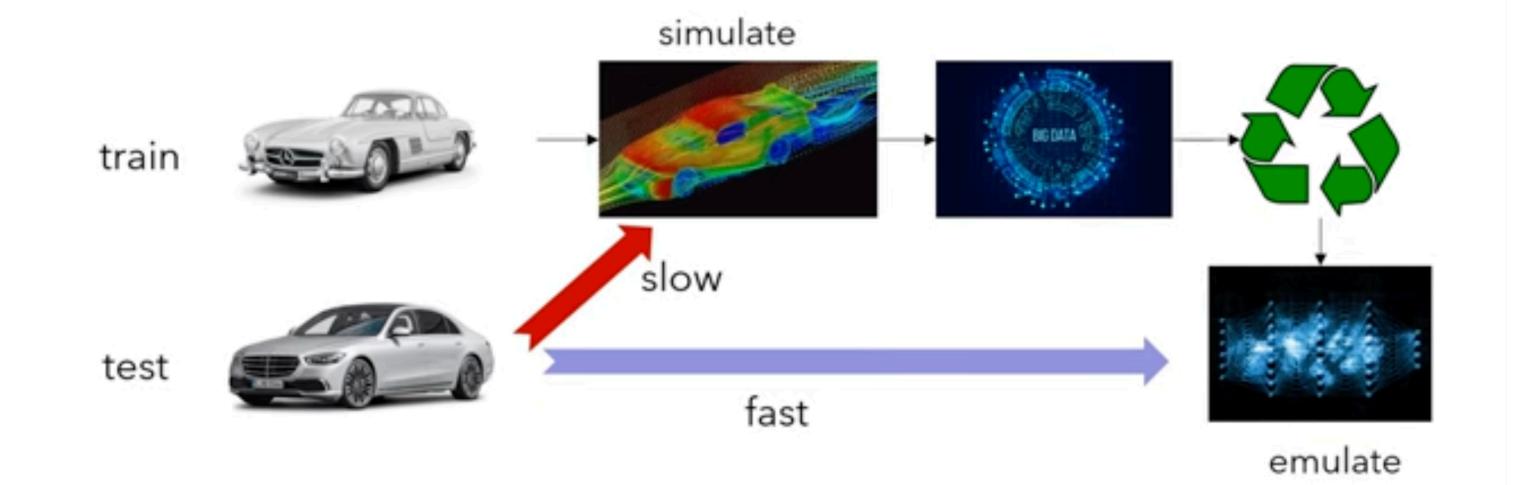
"The underlying physical laws necessary for the mathematical theory of a large part of physics and the whole of chemistry are thus completely known, and the difficulty is only that the exact application of these laws leads to equations much too complicated to be soluble."

-PAUL DIRAC

Simulation & Emulators

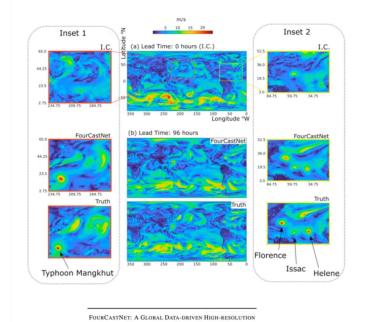
Amortization through Simulation

Simulate train NN surrogate emulate

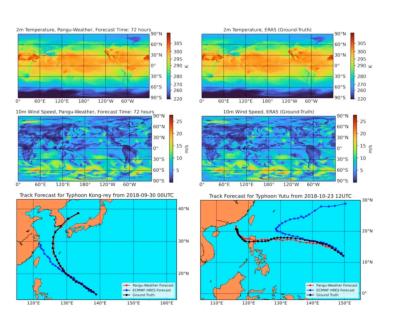




NN predicts the Weather 10,000 times faster

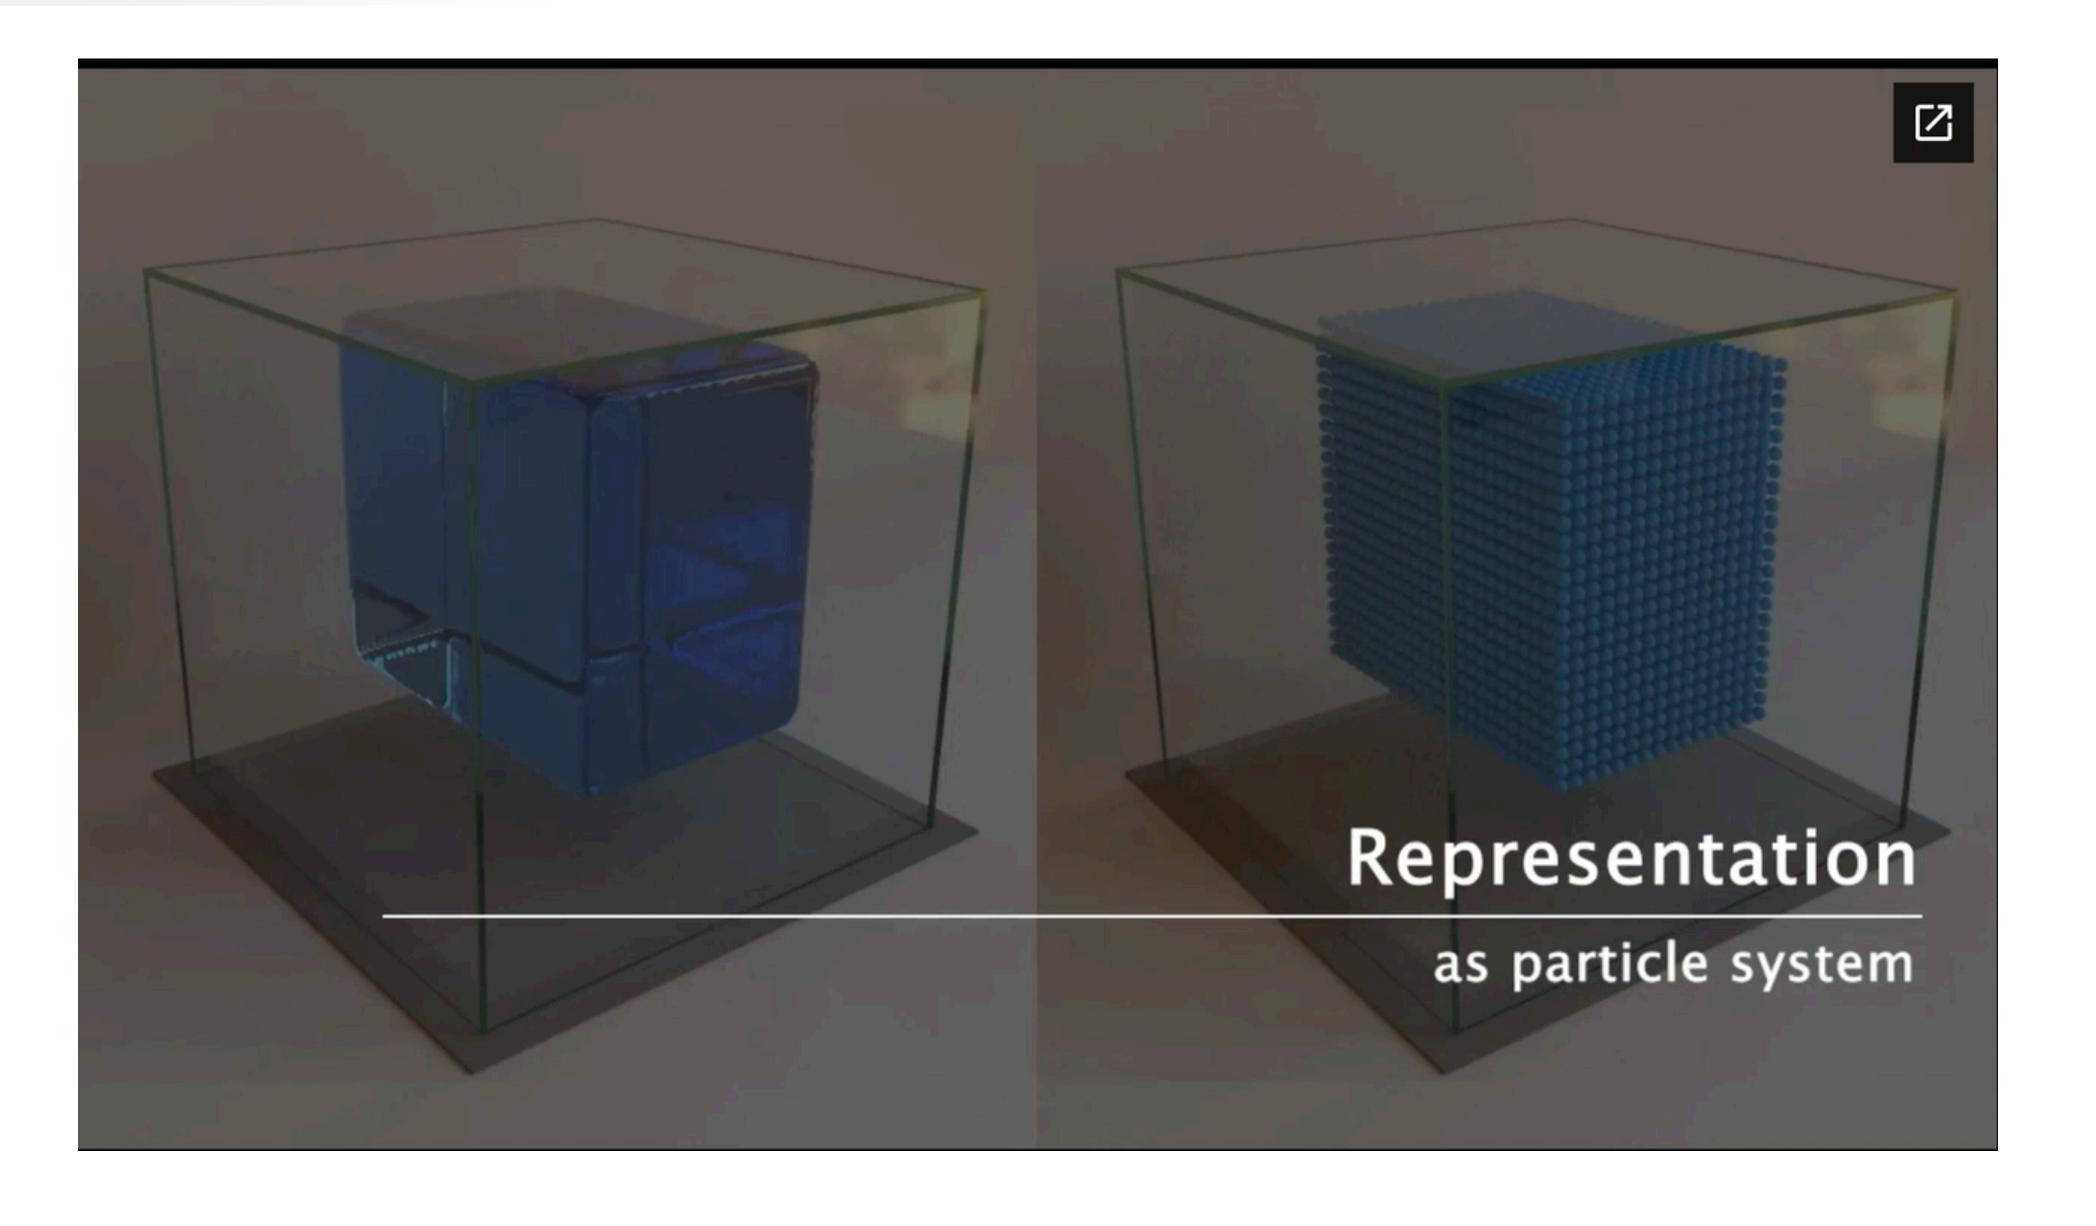




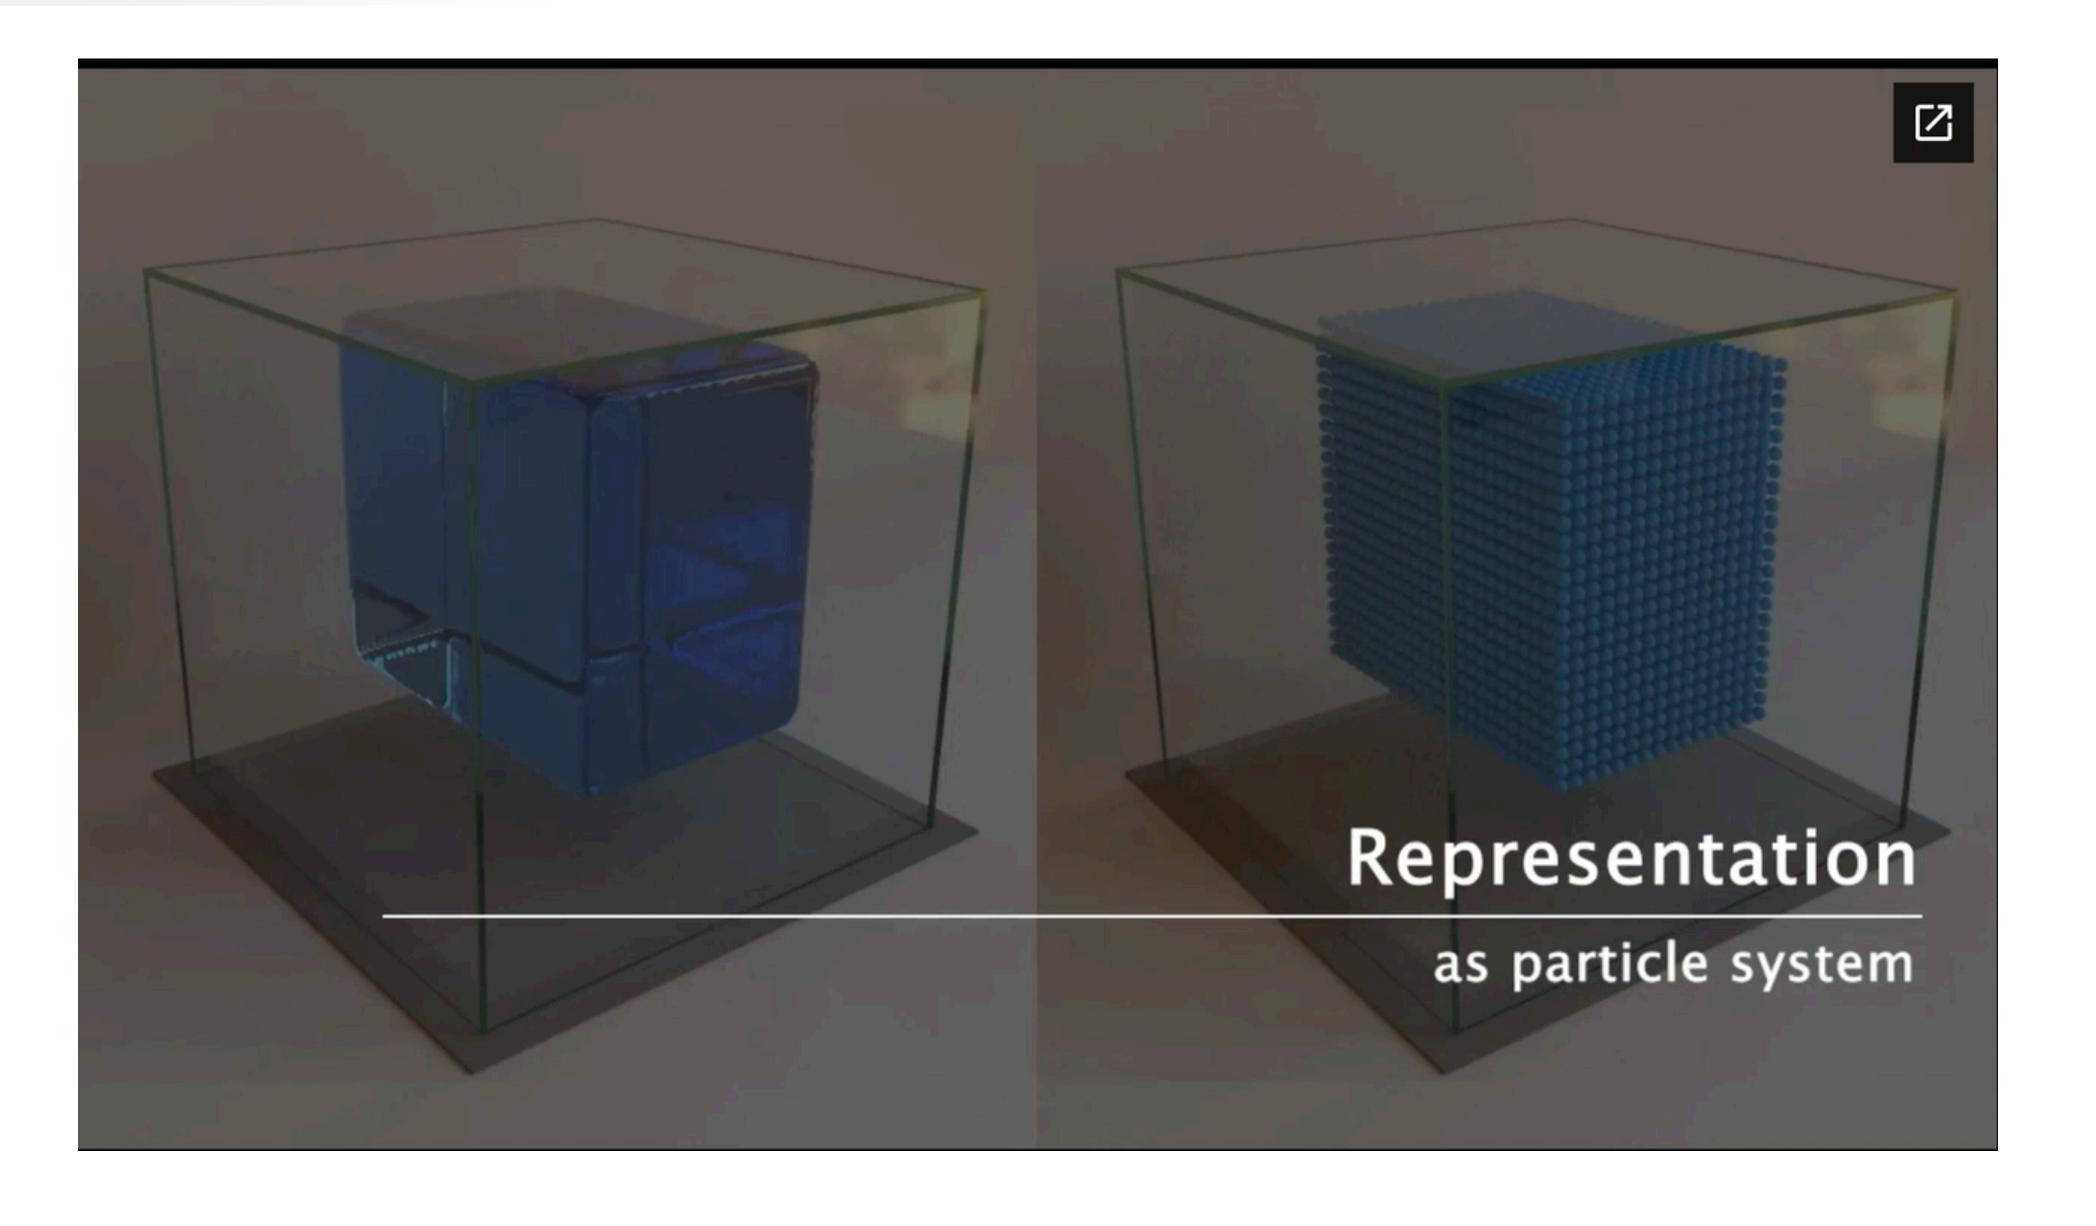


Pangu-Weather: A 3D High-Resolution System for Fast and Accurate Global Weather Forecast Kaifeng Bi, Lingxi Xie, Hengheng Zhang, Xin Chen, Xiaotao Gu, and Qi Tian^{FS}, Fellow, IEEE

AI/ML as emulators of complex simulations



AI/ML as emulators of complex simulations



In fusion energy



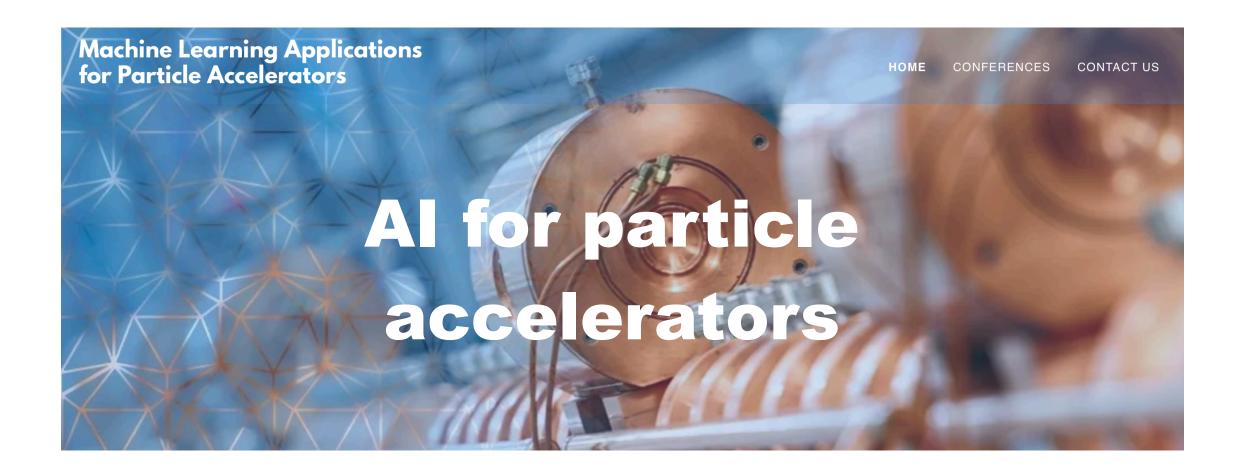
Accelerating fusion science through learned plasma control

February 16, 2022



Successfully controlling the nuclear fusion plasma in a tokamak with deep reinforcement learning

Verena Kain's talk



AI for particle accelerators, EuCAIF, V. Kain, 01-May-2024

RL4AA - workshop



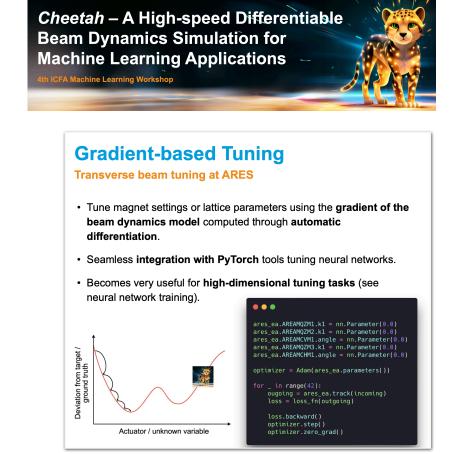
Pushing the frontiers of RL for accelerators \rightarrow autonomous accelerators.



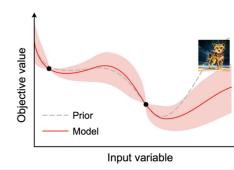
Differentiable simulation codes



Optimisation algorithms work best and are most sample-efficient with gradient information of the objective function.



- A physics-informed prior can help improve the performance of BO by preventing over-exploitation.
- Cheetah's differentiability allows efficient acquisition function optimisation using gradient descent methods in modern BO packages like BoTorch.
- Has well-defined behaviour and does not need data to train like neural network priors.
- Can be used in combination with gradient-based system identification to overcome model inaccuracies.

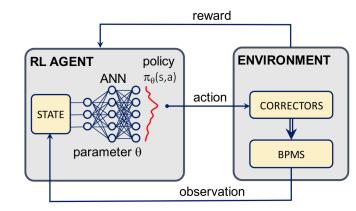


AI for particle accelerators, EuCAIF, V. Kain, 01-May-2024

Reinforcement Learning (RL)



Learn dynamics (once and for all) through trial-and-error, no exploration after training!



RL setup for trajectory steering

Next generation accelerators to be built <u>for</u> RL:

→ fast executing (accurate) simulation / digital twin

for training

→ instrumentation designed with control algorithm

RL elegant (if not ideal) solution, but online training often not possible!

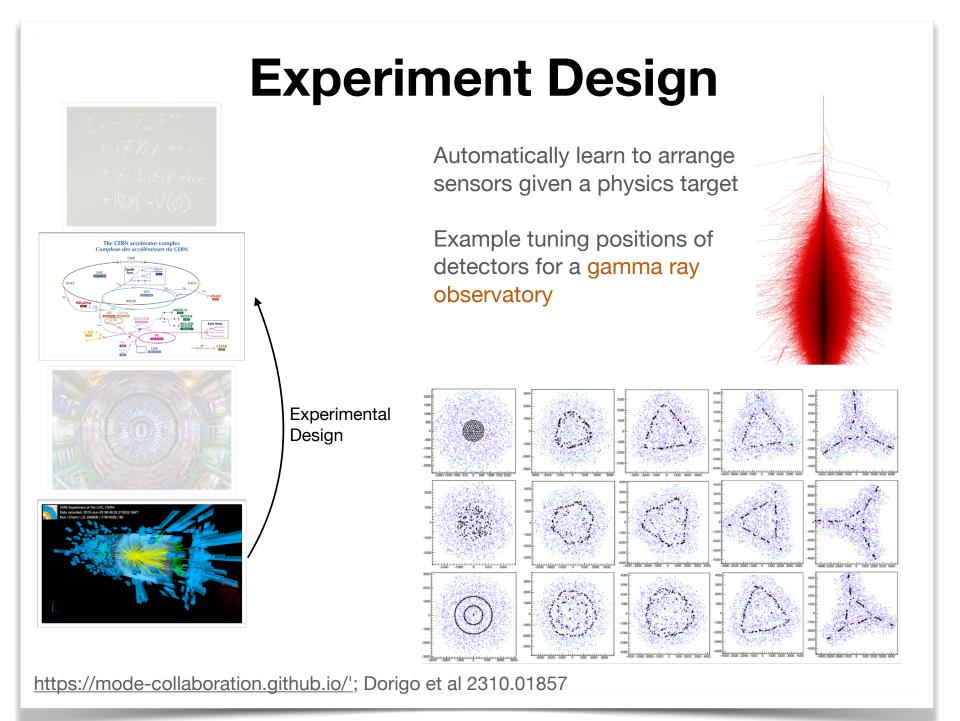
- Not sample-efficient enough
- Safety constraints
- → RL (like MPC) needs to be built into accelerator design.

AI for particle accelerators, EuCAIF, V. Kain, 01-May-2024

AI for particle accelerators, EuCAIF, V. Kain, 01-May-2024

Surrogates and differentiable programming





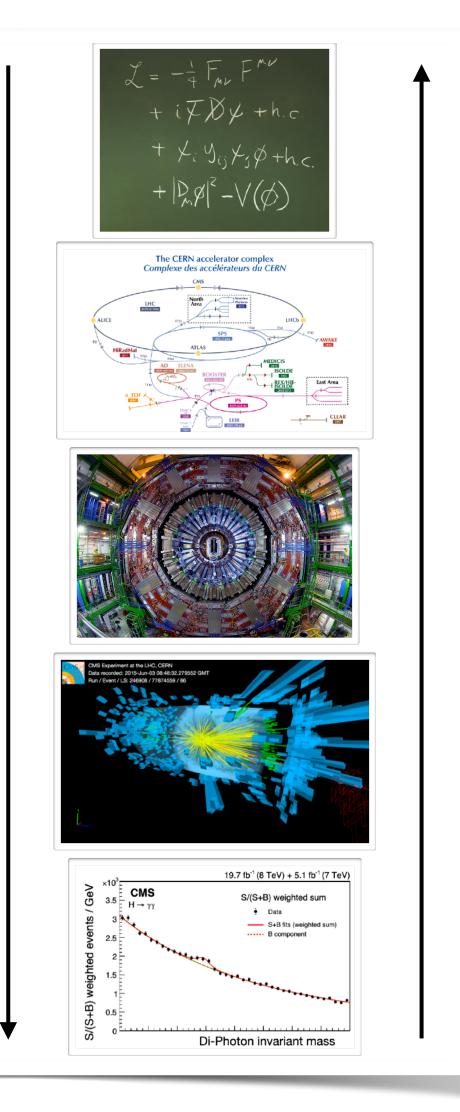
Differentiable versions of all steps in the particle physics processing chain

Either as ML-based surrogate models

Or via e.g. differentiable programming

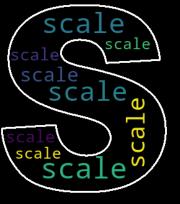
What can we do with this?

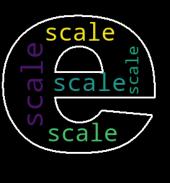
Heinrich, Kagan 2308.16680



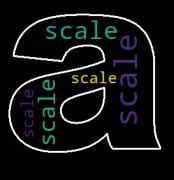
Slides from Gregor Kasieczka's talk

Inductive Bias Compositionality Relationships Symmetry Causality

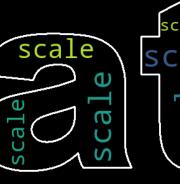
















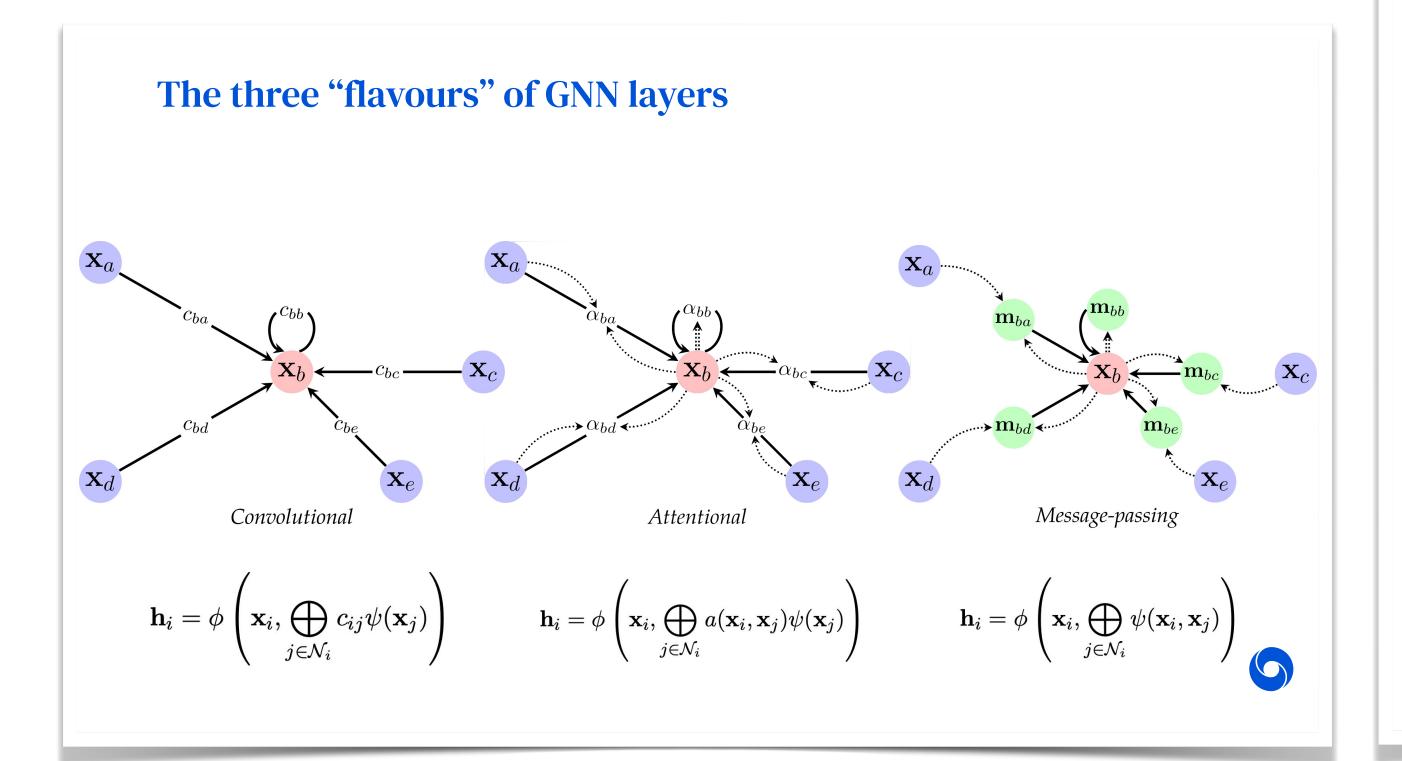


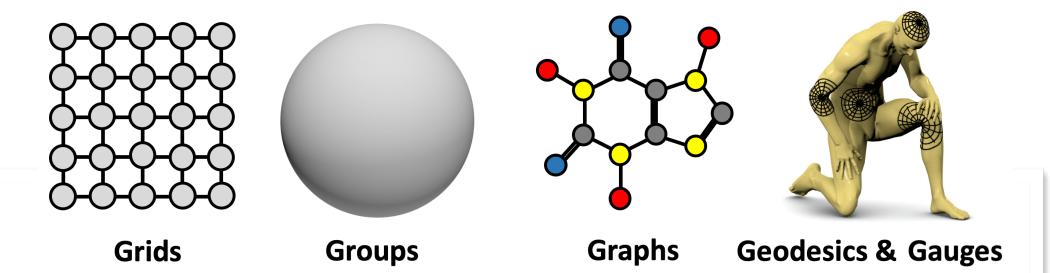


Geometric Deep Learning

A hot topic in AI/ML research

- How do we incorporate domain knowledge?
- Should we? "The bitter lesson"





Geometric Deep Learning Grids, Groups, Graphs, Geodesics, and Gauges

Michael M. Bronstein¹, Joan Bruna², Taco Cohen³, Petar Veličković⁴

May 4, 2021

2021

May

¹Imperial College London / USI IDSIA / Twitter

²New York University

³Qualcomm AI Research. Qualcomm AI Research is an initiative of Qualcomm Technologies, Inc.

⁴DeepMind

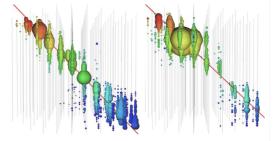
Particle physics & astrophysics are highlighted in review

 Continues to be an effective connection with Al community

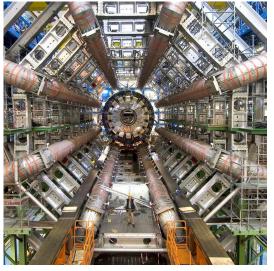
6. PROBLEMS AND APPLICATIONS | 113

of particular interest, since neutrinos interact only very rarely with matter, and thus travel enormous distances practically unaffected. Detecting neutrinos allows to observe objects inaccessible to optical telescopes, but requires enormously-sized detectors – the IceCube neutrino observatory uses a cubic kilometer of Antarctic ice shelf on the South Pole as its detector. Detecting high-energy neutrinos can possibly shed lights on some of the most mysterious objects in the Universe, such as blazars and black holes. Choma et al. (2018) used a Geometric neural network to model the irregular geometry of the IceCube neutrino detector, showing significantly better performance in detecting neutrinos coming from astrophysical sources and separating them from background events.

While neutrino astronomy offers a big promise in the study of the Cosmos, traditional optical and radio telescopes are still the 'battle horses' of astronomers. With these traditional instruments, Geometric Deep Learning can still offer new methodologies for data analysis. For example, Scaife and Porter (2021) used rotationally-equivariant CNNs for the classification of radio galaxies, and McEwen et al. (2021) used spherical CNNs for the analysis of cosmic microwave background radiation, a relic from the Big Bang that might shed light on the formation of the primordial Universe. As we already mentioned, such signals are naturally represented on the sphere and equivariant neural networks are an appropriate tool to study them.



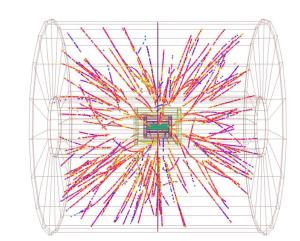
The characteristic pattern of light deposition in IceCube detector from background events (muon bundles, left) and astrophysical neutrinos (high-energy single muon, right). Choma et al. (2018)



Part of the Large Hadron Collider detectors.

Particle physics and astrophysics High energy physicists were perhaps among the first domain experts in the field of natural sciences to embrace the new shiny tool, graph neural networks. In a recent review paper, Shlomi et al. (2020) note that machine learning has historically been heavily used in particle physics experiments, either to learn complicated inverse functions allowing to infer the underlying physics process from the information measured in the detector, or to perform classification and regression tasks. For the latter, it was often necessary to force the data into an unnatural representation such as grid, in order to be able to used standard deep learning architectures such as CNN. Yet, many problems in physics involve data in the form of unordered sets with rich relations and interactions, which can be naturally represented as graphs.

One important application in high-energy physics is the reconstruction and classification of *particle jets* – sprays of stable particles arising from multiple successive interaction and decays of particles originating from a single initial event. In the Large Hardon Collider, the largest and best-known particle accelerator built at CERN, such jet are the result of collisions of protons at nearly the speed of light. These collisions produce massive particles, such as the long though-for Higgs boson or the top quark. The identification and classification of collision events is of crucial importance, as it might provide experimental evidence to the existence of new particles.



Example of a particle jet

Multiple Geometric Deep Learning approaches have recently been proposed for particle jet classification task, e.g. by Komiske et al. (2019) and Qu and Gouskos (2019), based on DeepSet and Dynamic Graph CNN architectures, respectively. More recently, there has also been interest in developing specialsed architectures derived from physics consideration and incorporating inductive biases consistent with Hamiltonian or Lagrangian mechanics (see e.g. Sanchez-Gonzalez et al. (2019); Cranmer et al. (2020)), equivariant to the Lorentz group (a fundamental symmetry of space and time in physics) (Bogatskiy et al., 2020), or even incorporating symbolic reasoning (Cranmer et al., 2019) and capable of learning physical laws from data. Such approaches are more interpretable (and thus considered more 'trustworthy' by domain experts) and also offer better generalisation.

Besides particle accelerators, particle detectors are now being used by astrophysicist for *multi-messenger astronomy* – a new way of coordinated observation of disparate signals, such as electromagnetic radiation, gravitational waves, and neutrinos, coming from the same source. Neutrino astronomy is

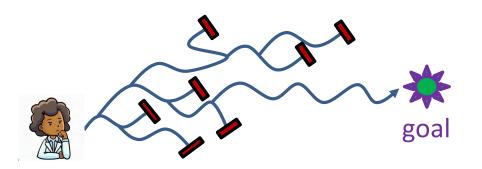
Scientific Understanding

Matt Schwartz's talk

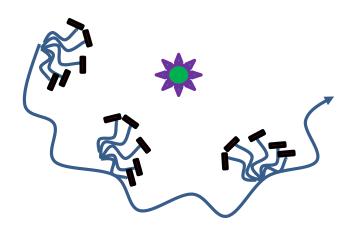


Theoretical High Energy Physics may have stalled

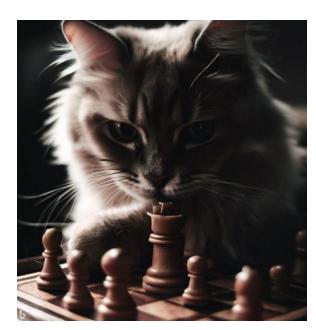
In the past, we made progress depsite many dead ends



Are we even making forward progress anymore?



Maybe the problems are just too difficult (for us)



Could a cat ever learn to play chess?

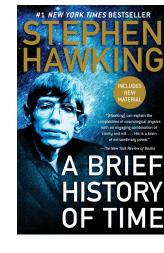
• Humans have limits too

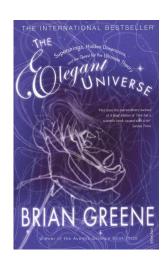
Beyond augmented intelligence

Suppose a machine understands the theory of everything but we don't

- e.g. can calculate the fine-structure constant from scratch
- e.g. can preduct the endpoint of black-hole evaporation

Is this enough or do we need to understand it too?





• The authors of **Popular science books** understand the details; we just get the general idea

I don't understand the proof of Fermat's last theorem

- I'm glad that somebody does
- Does it matter that the person is human?

If a machine understands fundamental physics it can

- 1. Dumb it down so we can get the general idea
- 2. Find practical applications

Is this what we want? No.
But maybe it's the best we will get.

Because of AI, I am now optimistic for substantive progress in high-energy theory in my lifetime

Scientific understanding

How does Al enable or enhance scientific understanding?

On scientific understanding with artificial intelligence

Mario Krenn, ^{1,2,3,4,*} Robert Pollice, ^{2,3} Si Yue Guo, ² Matteo Aldeghi, ^{2,3,4} Alba Cervera-Lierta, ^{2,3} Pascal Friederich, ^{2,3,5} Gabriel dos Passos Gomes, ^{2,3} Florian Häse, ^{2,3,4,6} Adrian Jinich, ⁷ AkshatKumar Nigam, ^{2,3} Zhenpeng Yao, ^{2,8,9,10} and Alán Aspuru-Guzik^{2,3,4,11,†}

Scientific understanding

How does Al enable or enhance scientific understanding?

On scientific understanding with artificial intelligence

Mario Krenn, ^{1,2,3,4,*} Robert Pollice, ^{2,3} Si Yue Guo, ² Matteo Aldeghi, ^{2,3,4} Alba Cervera-Lierta, ^{2,3} Pascal Friederich, ^{2,3,5} Gabriel dos Passos Gomes, ^{2,3} Florian Häse, ^{2,3,4,6} Adrian Jinich, ⁷ Akshat Kumar Nigam, ^{2,3} Zhenpeng Yao, ^{2,8,9,10} and Alán Aspuru-Guzik ^{2,3,4,11,†}

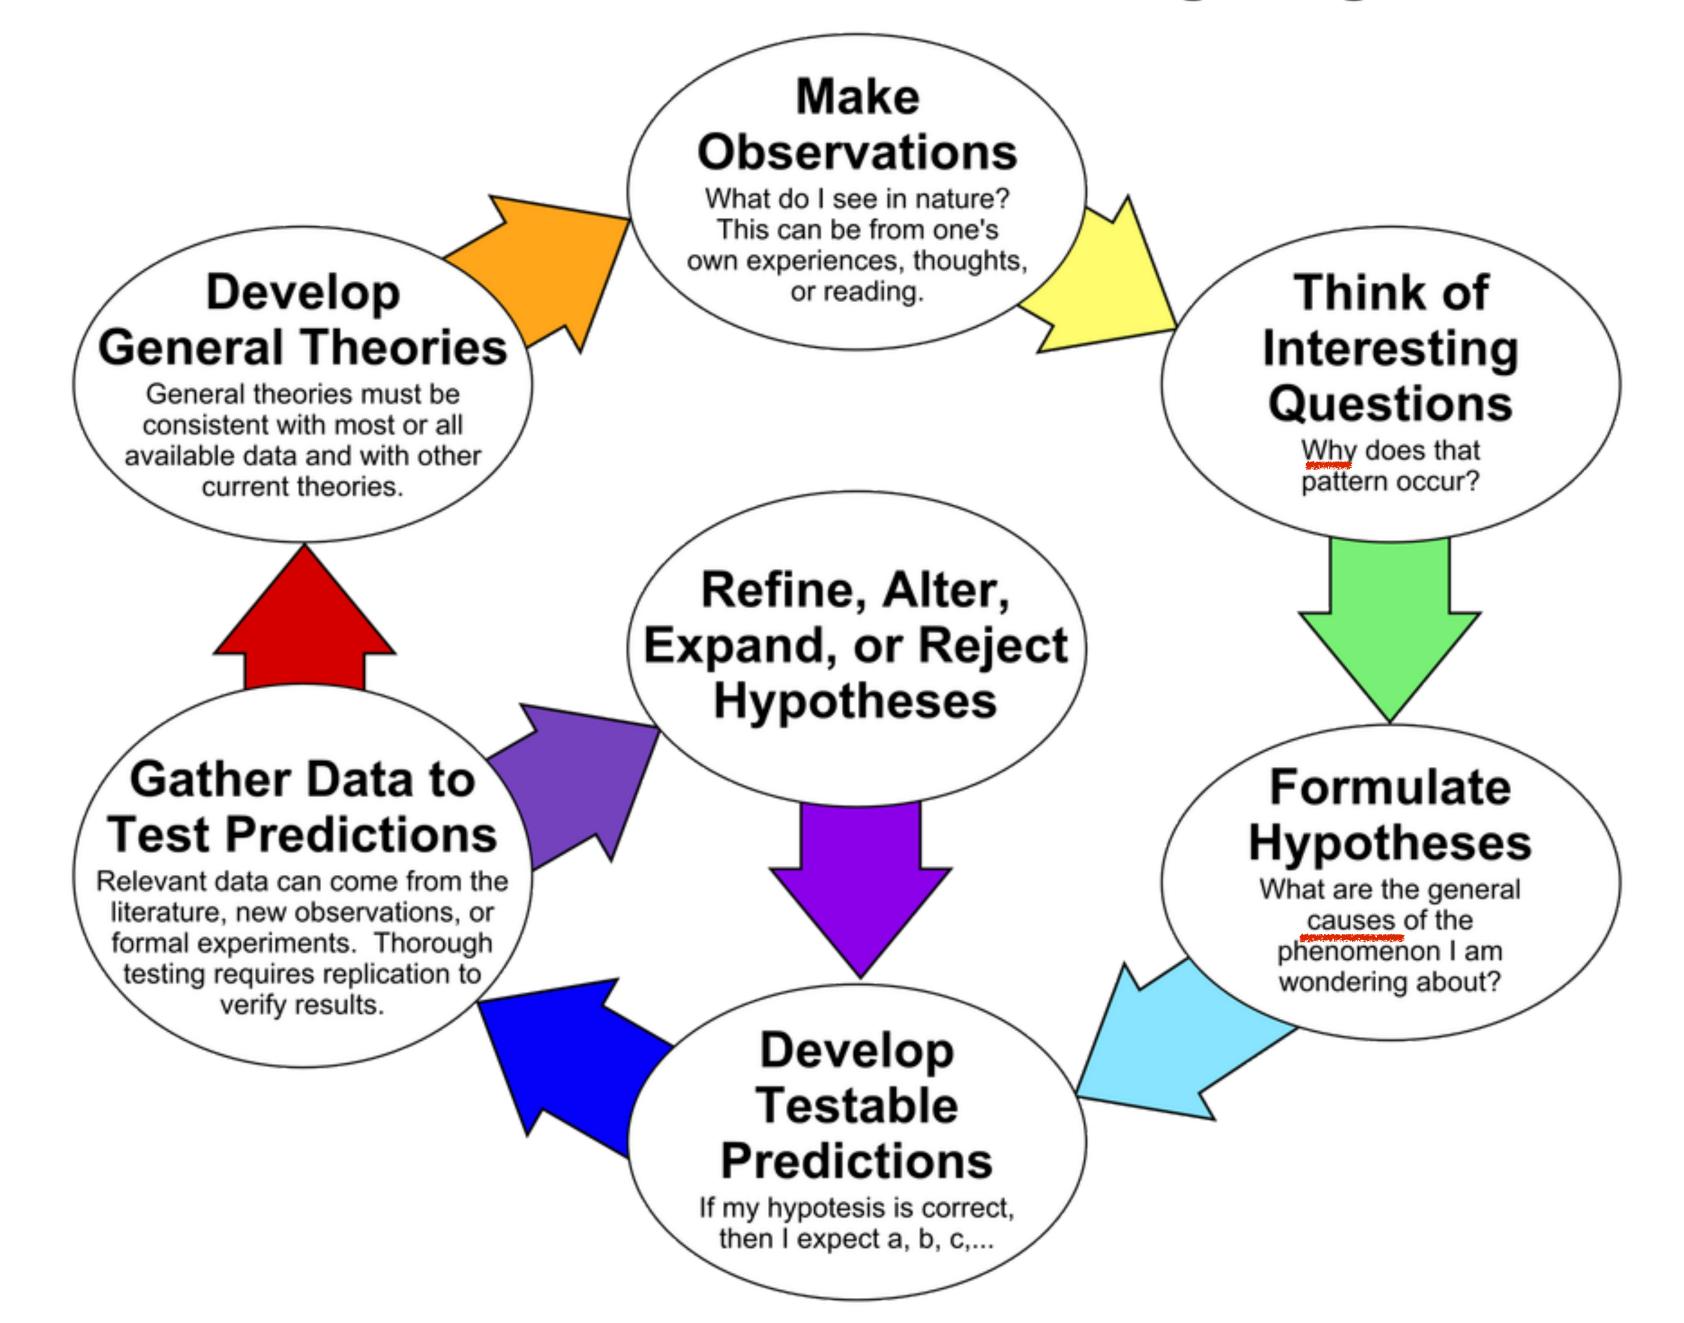
Scientific understanding

How does Al enable or enhance scientific understanding?

On scientific understanding with artificial intelligence

Mario Krenn, ^{1,2,3,4,*} Robert Pollice, ^{2,3} Si Yue Guo, ² Matteo Aldeghi, ^{2,3,4} Alba Cervera-Lierta, ^{2,3} Pascal Friederich, ^{2,3,5} Gabriel dos Passos Gomes, ^{2,3} Florian Häse, ^{2,3,4,6} Adrian Jinich, ⁷ Akshat Kumar Nigam, ^{2,3} Zhenpeng Yao, ^{2,8,9,10} and Alán Aspuru-Guzik ^{2,3,4,11,†}

The Scientific Method as an Ongoing Process



Looking forward

How does Al enable or enhance scientific understanding?

On scientific understanding with artificial intelligence

Mario Krenn, ^{1,2,3,4,*} Robert Pollice, ^{2,3} Si Yue Guo, ² Matteo Aldeghi, ^{2,3,4} Alba Cervera-Lierta, ^{2,3} Pascal Friederich, ^{2,3,5} Gabriel dos Passos Gomes, ^{2,3} Florian Häse, ^{2,3,4,6} Adrian Jinich, ⁷ Akshat Kumar Nigam, ^{2,3} Zhenpeng Yao, ^{2,8,9,10} and Alán Aspuru-Guzik ^{2,3,4,11,†}

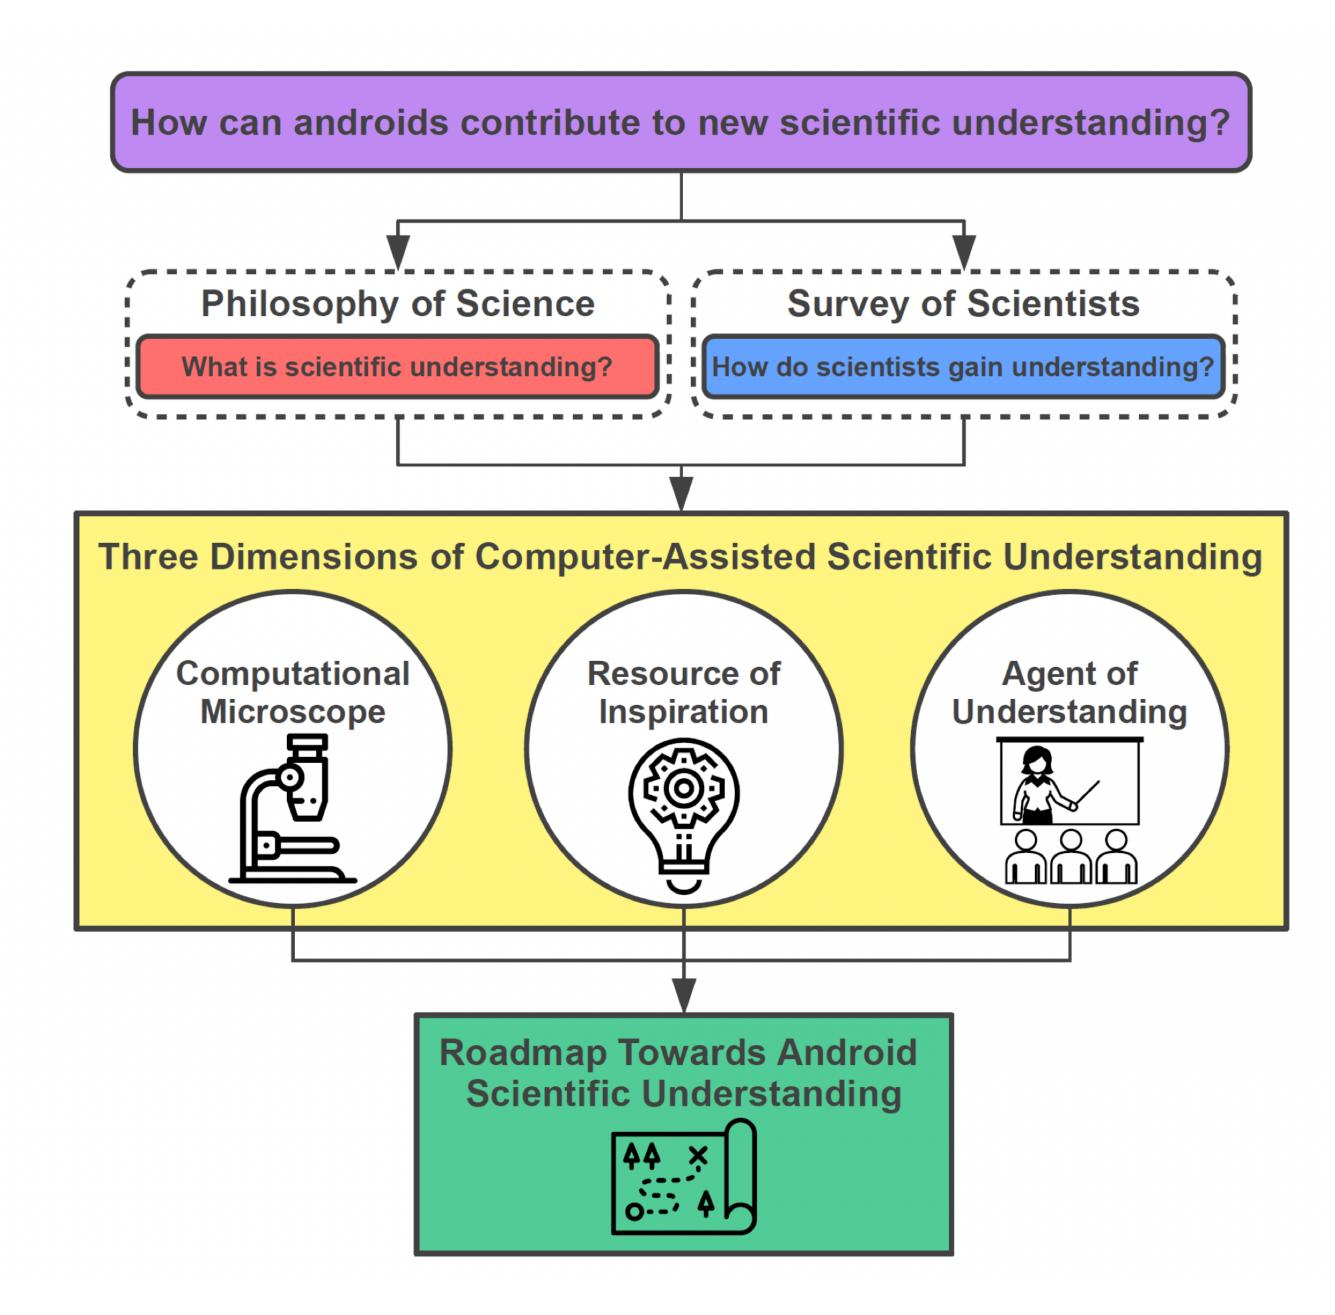
How does Al help?

An android can act

- I) as a computational microscope, providing information not (yet) attainable by experiment
- II) as a **resource of inspiration** or an artificial muse, expanding the scope of human imagination and creativity.

In those two classes, the human scientist is essential to take the new insight and inspiration and develop it to full understanding. Finally, an android can be

• III) an agent of understanding, replacing the human in generalizing observations and transferring scientific concepts to new phenomena.

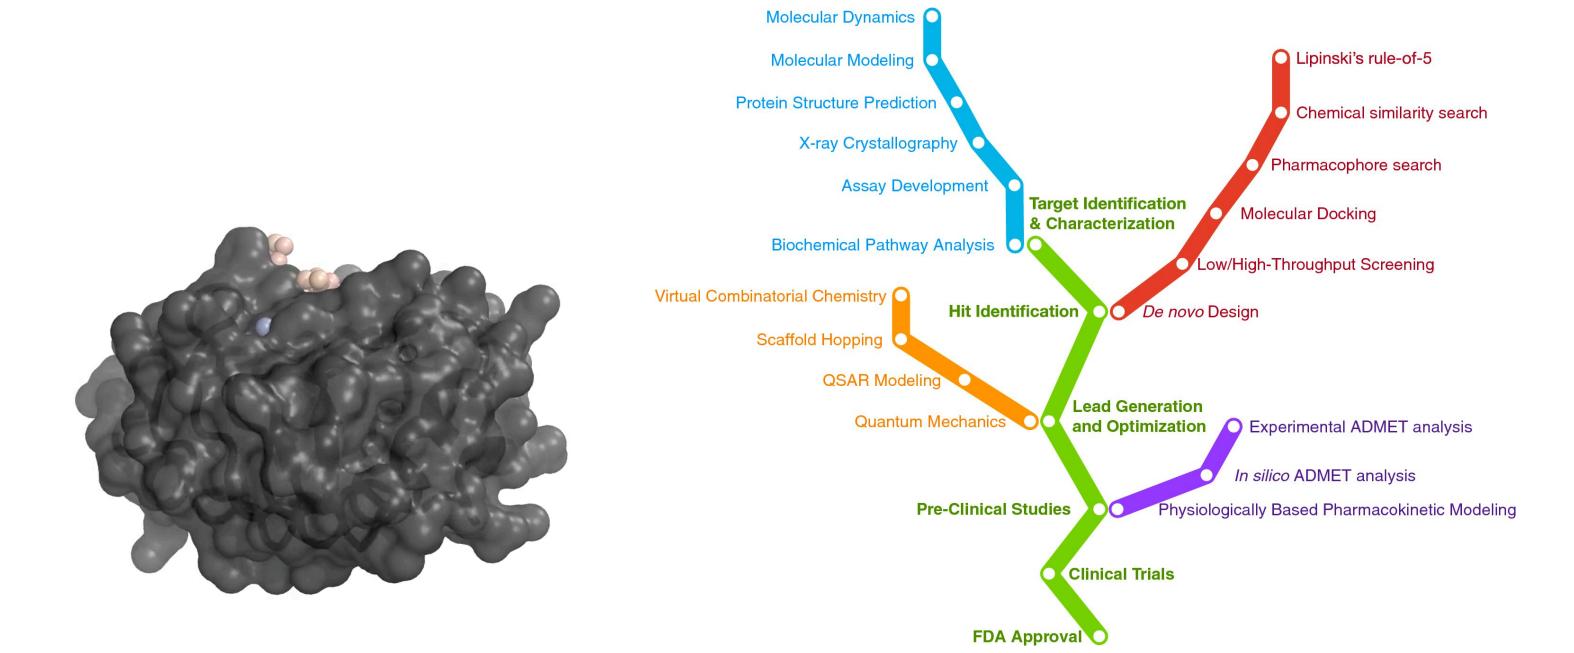


Use of ML in Physics vs. Molecules & Materials

Drug & Materials Discovery

Many uses of Al aimed at accelerating drug discovery and materials discovery

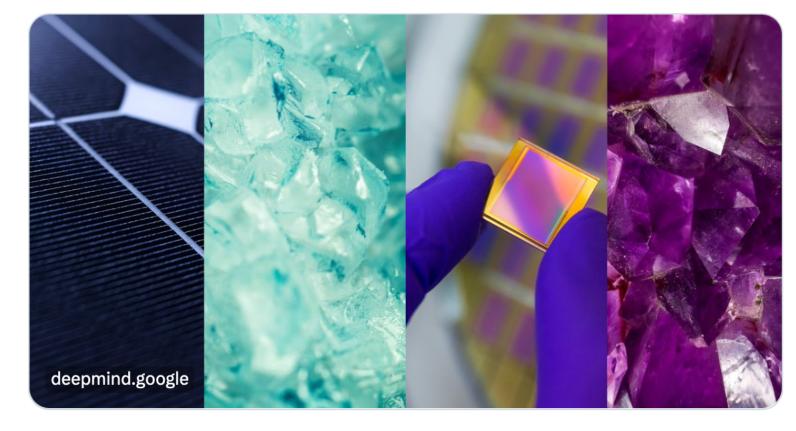
- Experimental follow up needed to confirm the predicted properties
- OK if the predictions are wrong as long as it accelerates the discovery process





We at @GoogleDeepMind are excited to announce #GNoME - an AI tool that has discovered 2.2 million new materials, and helps to predict material stability.

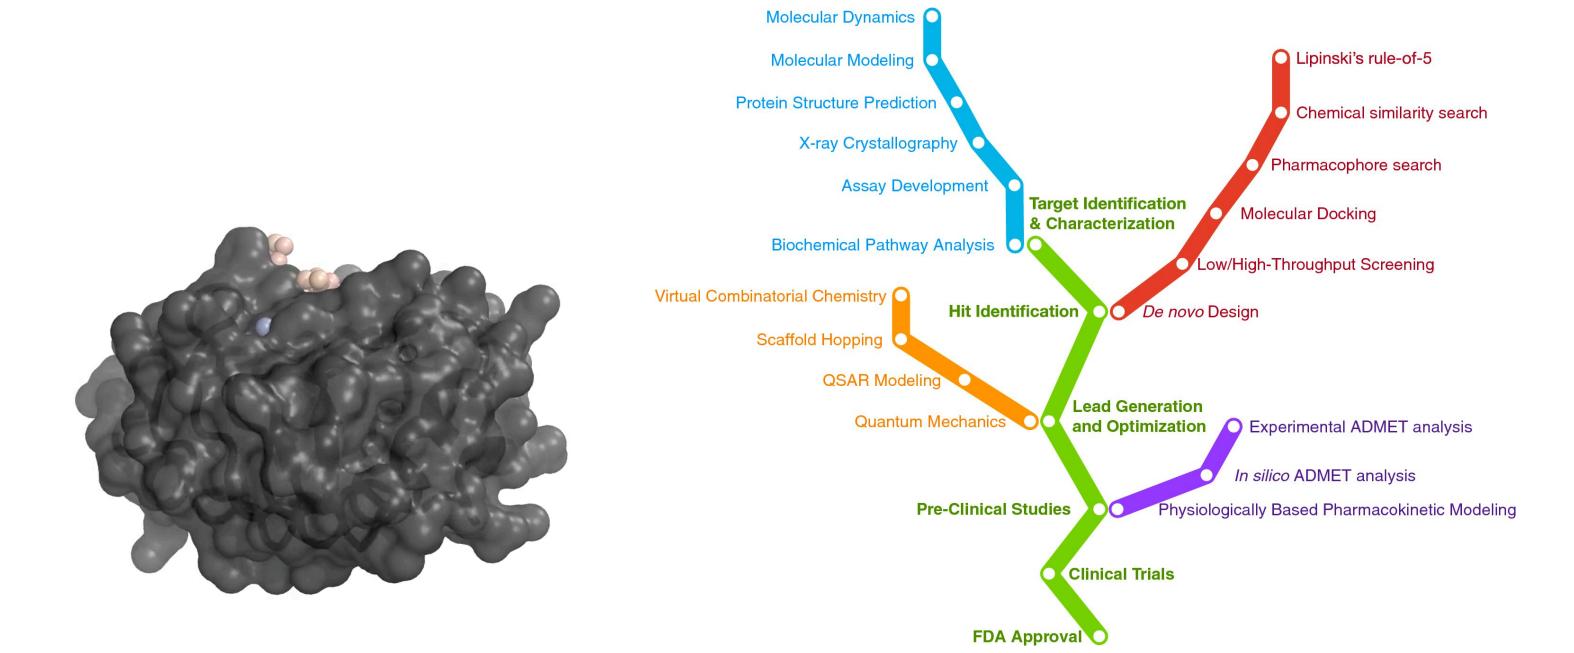
We're releasing 381K stable materials to help scientists pursue materials discovery breakthroughs.



Drug & Materials Discovery

Many uses of Al aimed at accelerating drug discovery and materials discovery

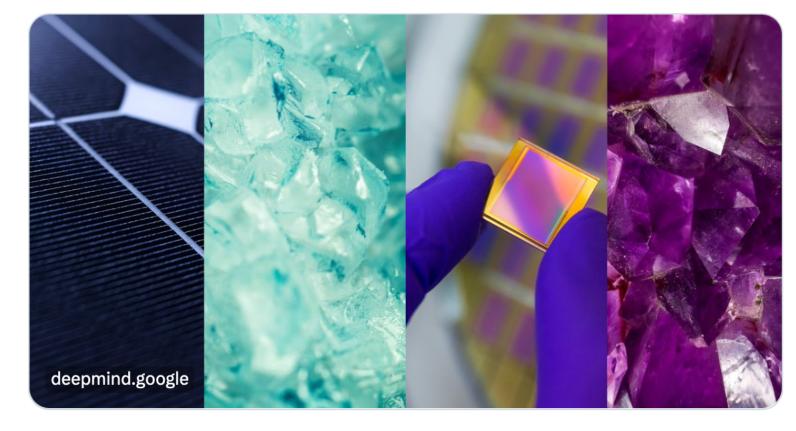
- Experimental follow up needed to confirm the predicted properties
- OK if the predictions are wrong as long as it accelerates the discovery process





We at @GoogleDeepMind are excited to announce #GNoME - an AI tool that has discovered 2.2 million new materials, and helps to predict material stability.

We're releasing 381K stable materials to help scientists pursue materials discovery breakthroughs.



Experimental Physics, Astrophysics, Cosmology

In contrast, AI/ML in experimental physics, astrophysics, and cosmology is often a component of a **hypothesis testing / statistical inference pipeline**.

- Robustness to distribution shift is important!
- Mistakes matter we need to be able to calibrate & perform uncertainty quantification!

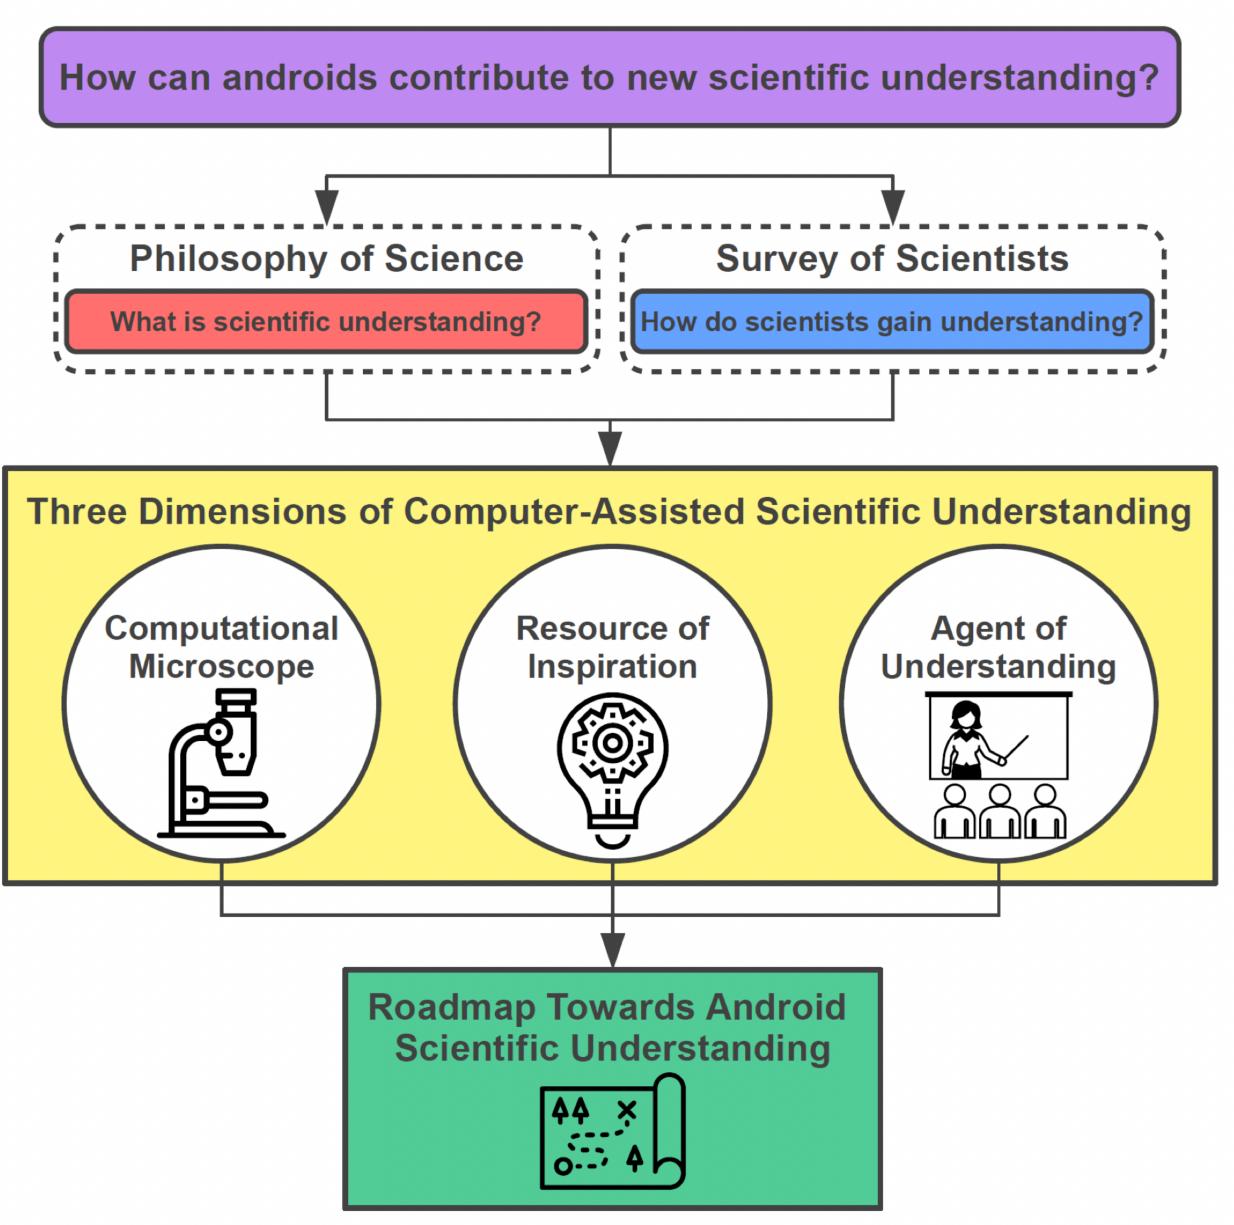


Data Collection

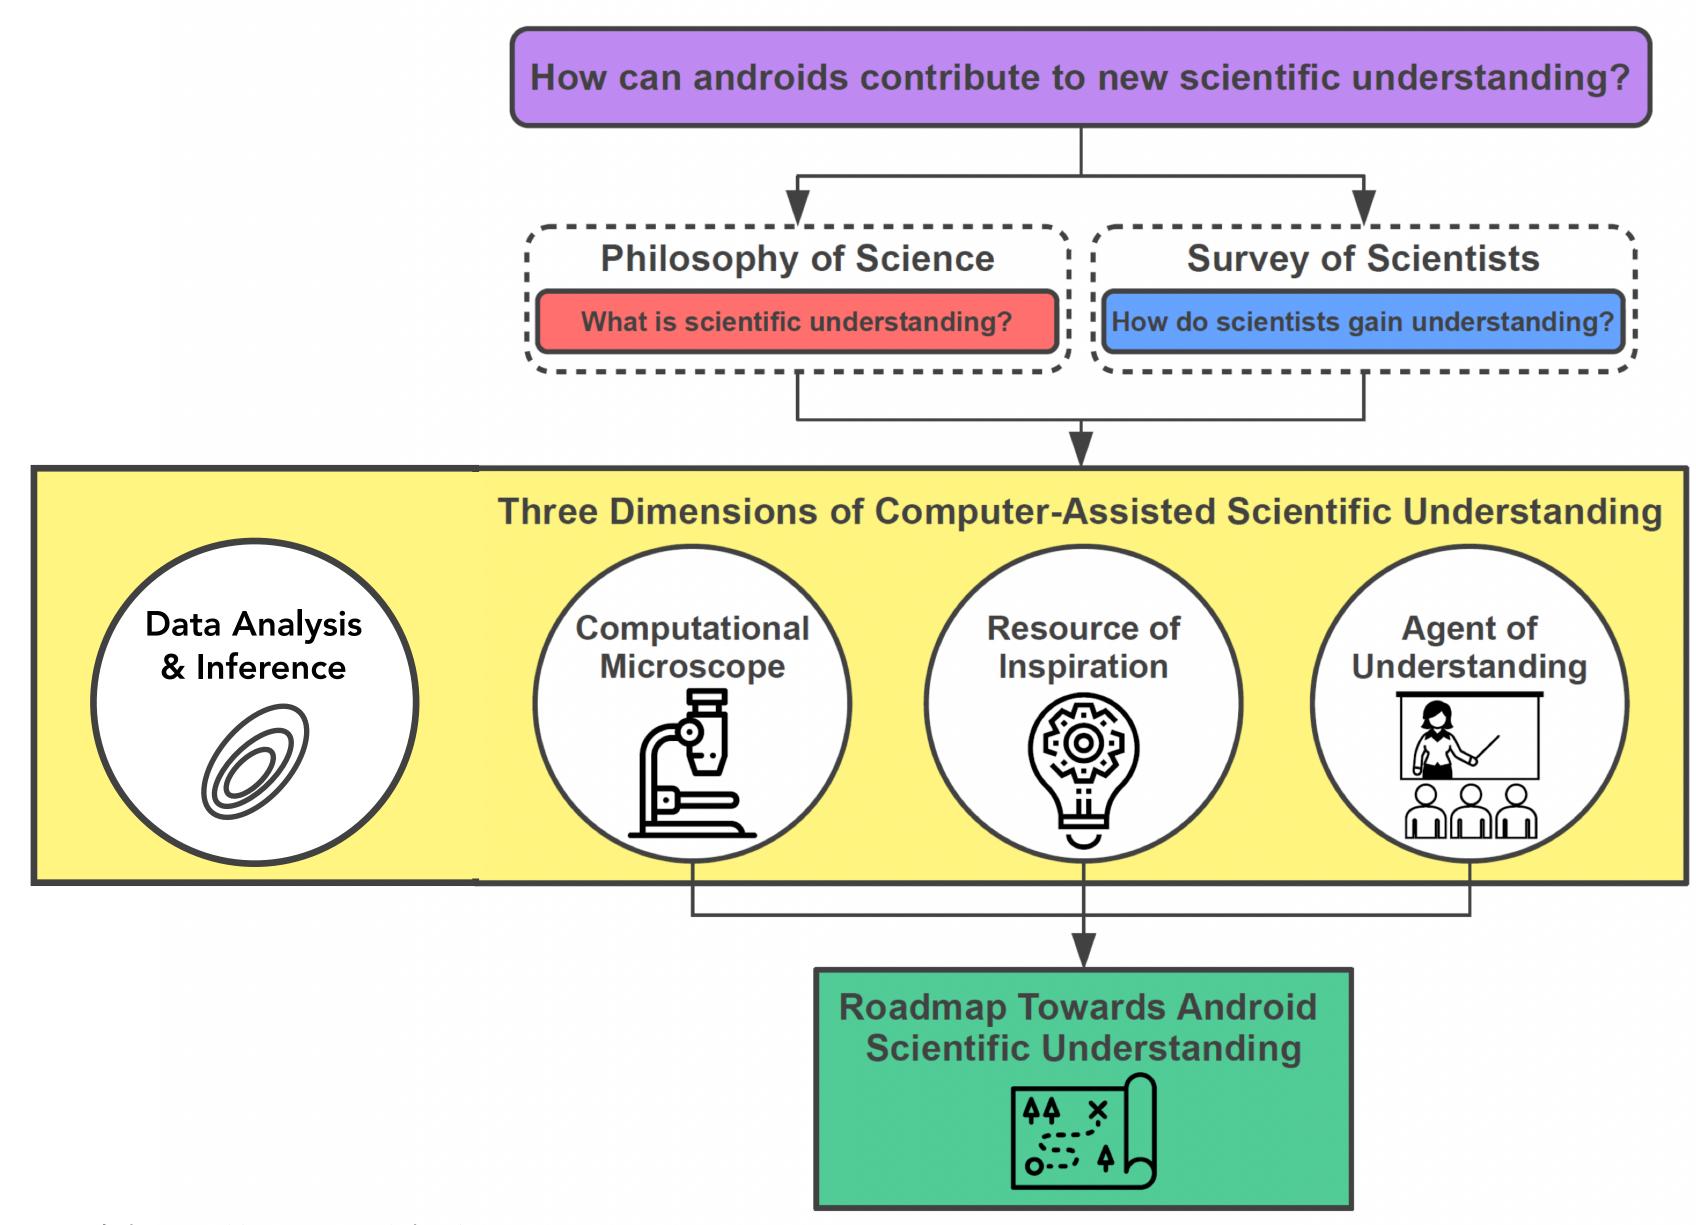
Data Analysis Pipeline

Scientific Claims

Computer-Assisted Scientific Understanding

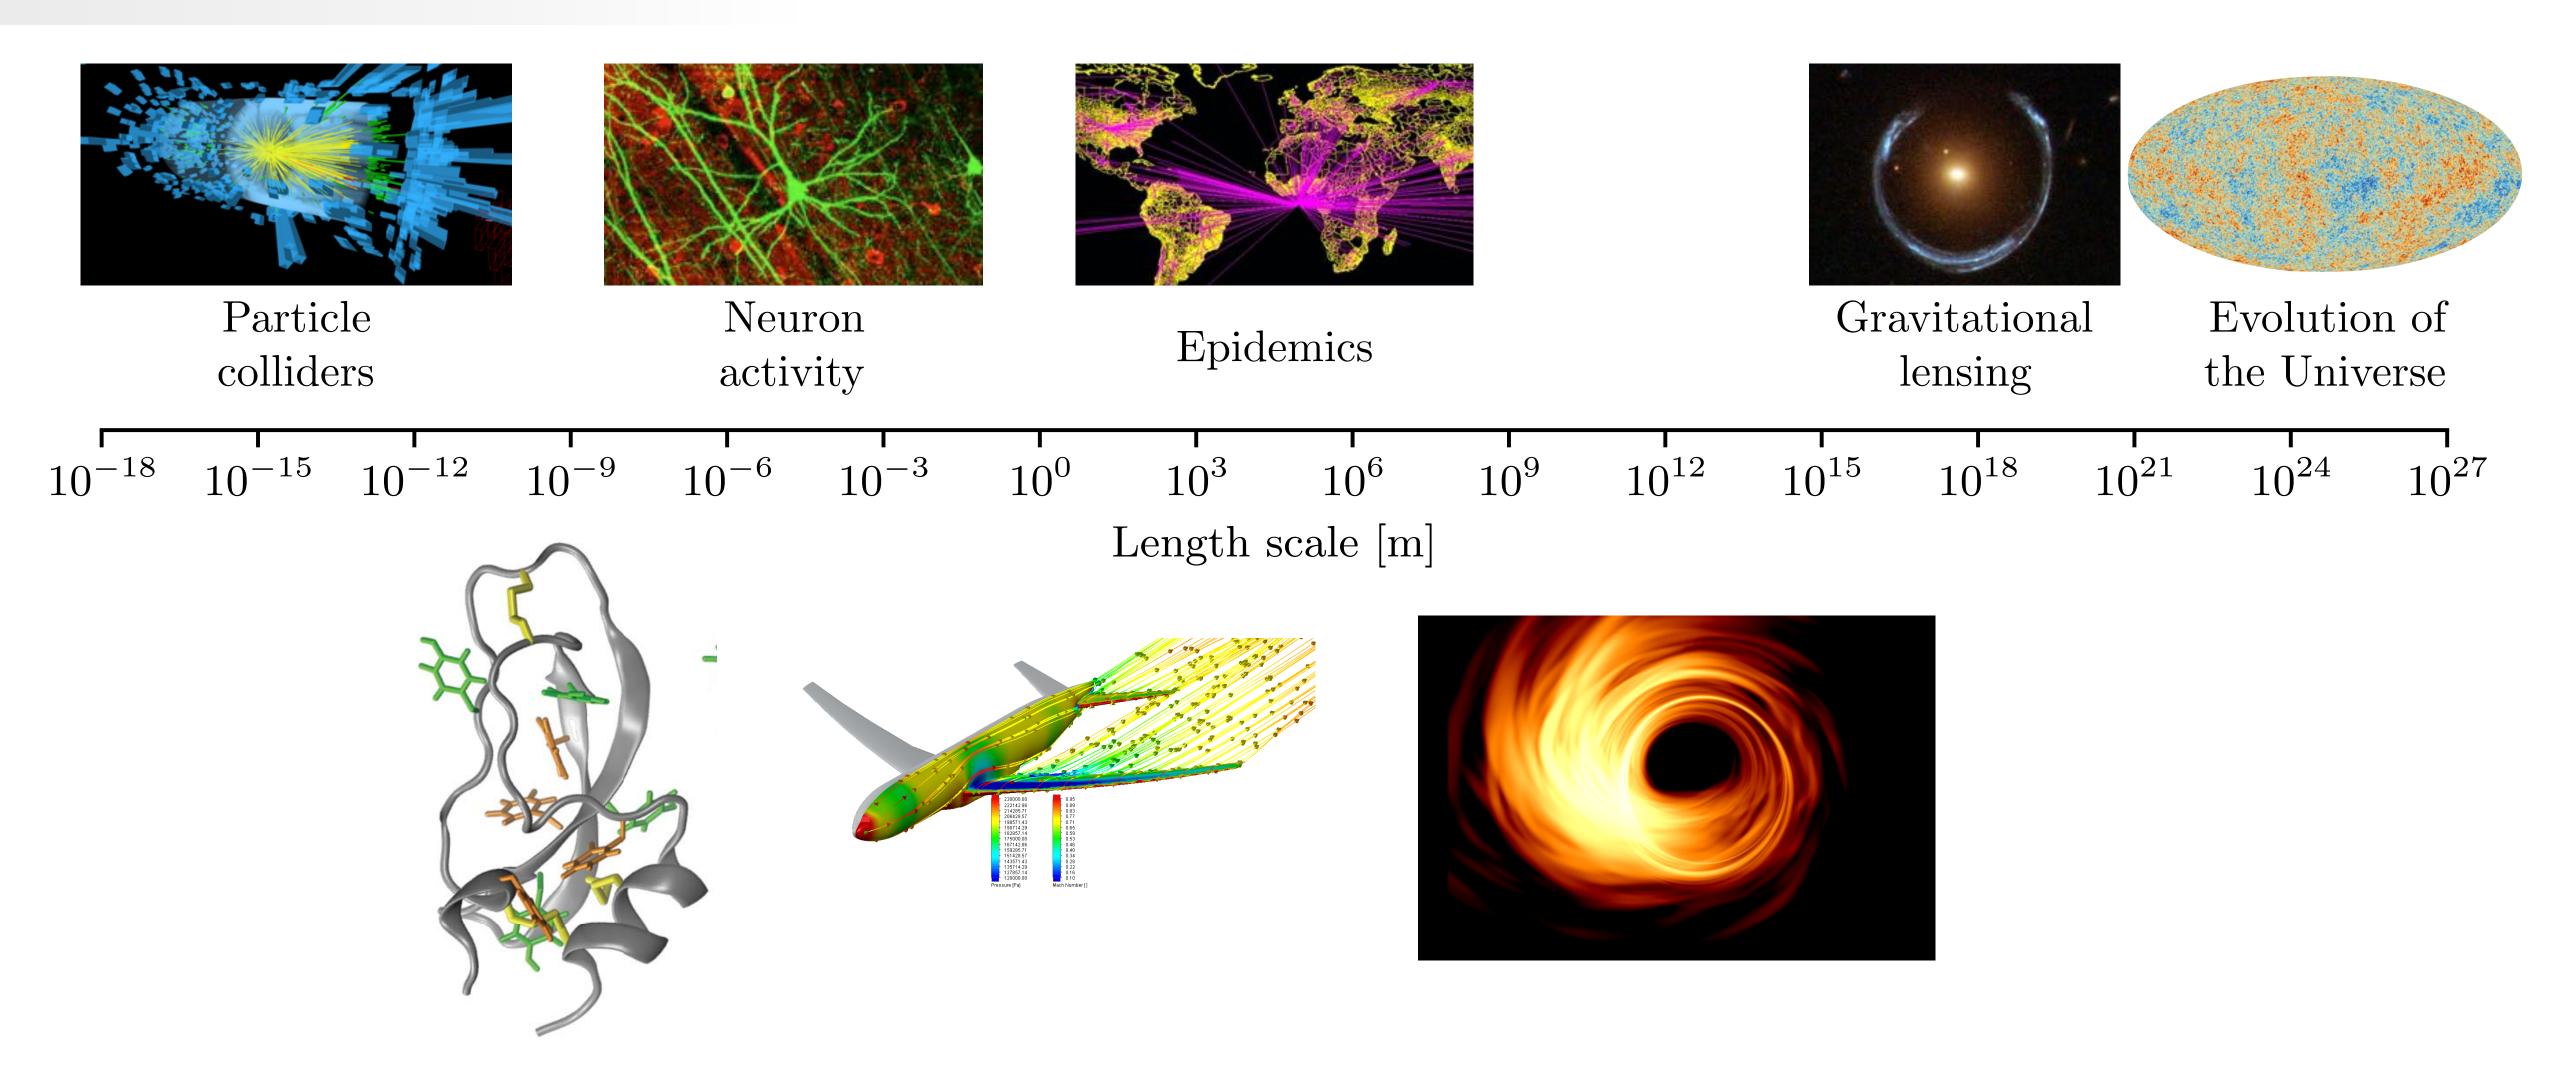


Computer-Assisted Scientific Understanding



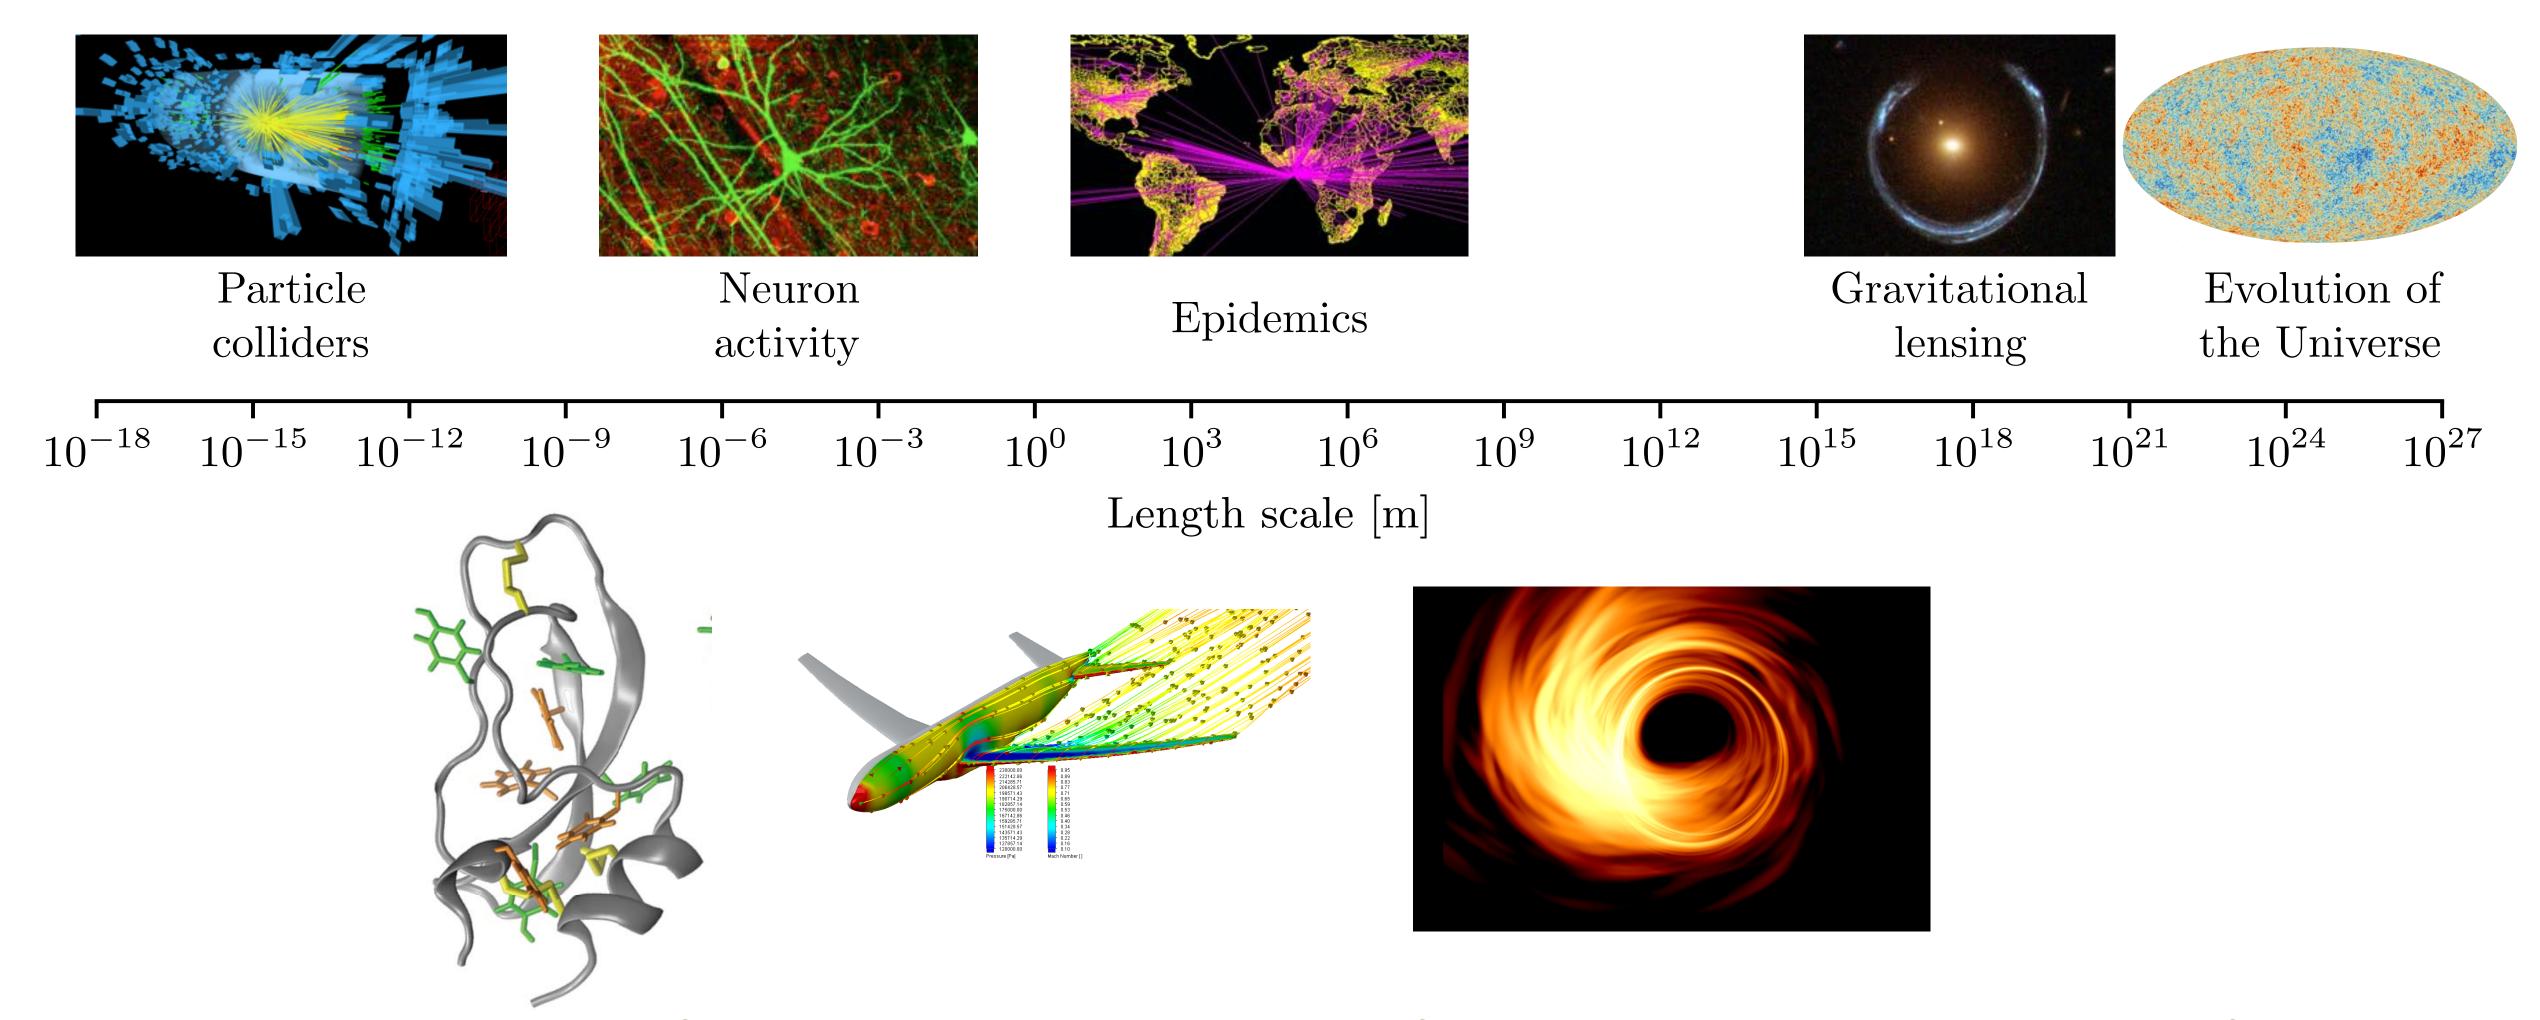
Simulation-based Inference

Science is replete with high-fidelity simulators



Simulators are causal, generative models of the data generating process

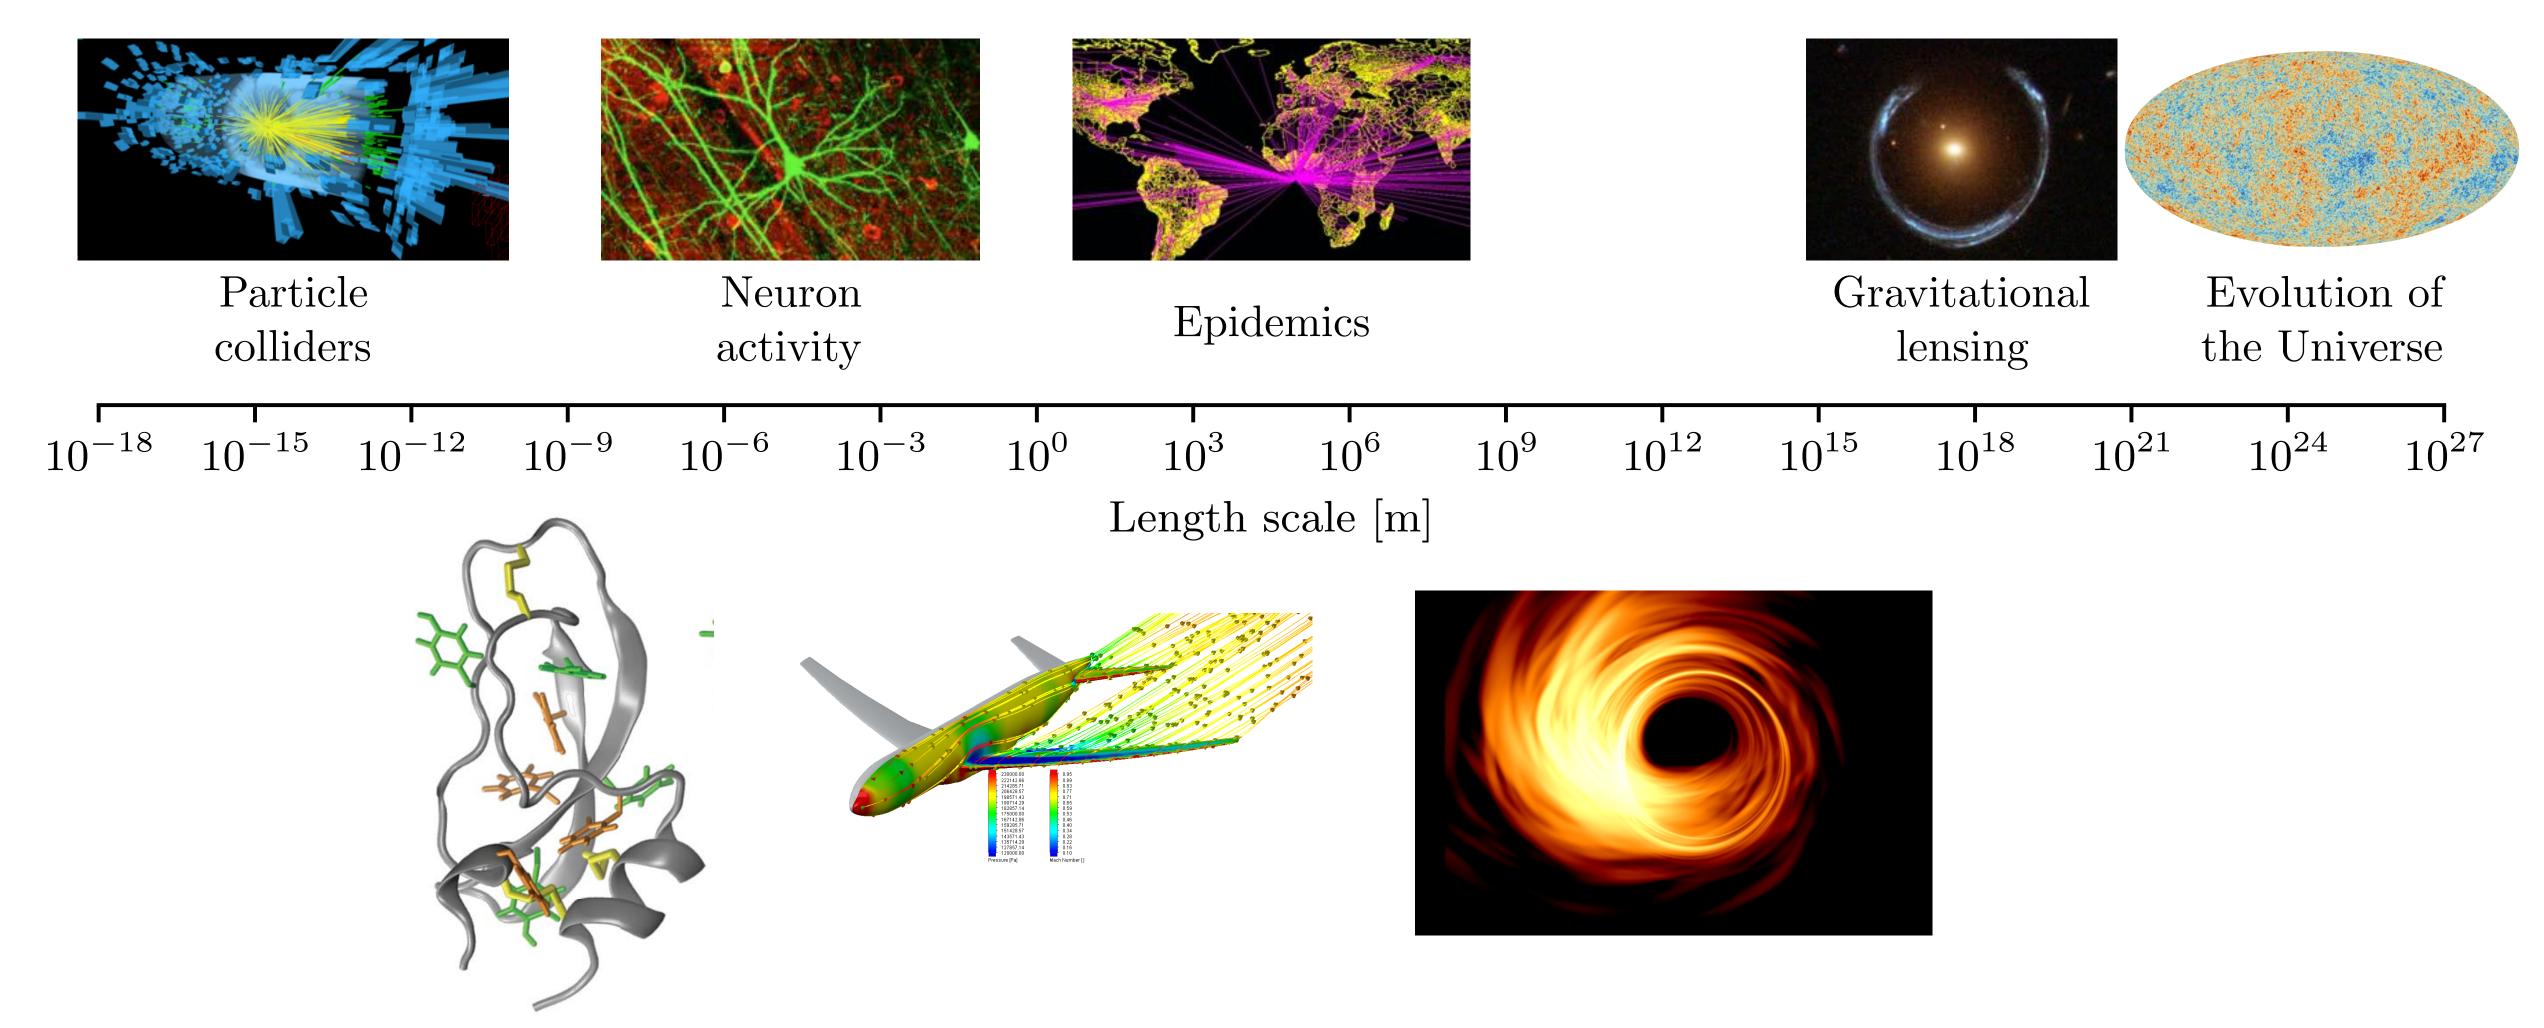
Science is replete with high-fidelity simulators



The expressiveness of programming languages facilitates the development of complex, high-fidelity simulations, and the power of modern computing provides the ability to generate synthetic data from them.

[Cranmer, Brehmer, Louppe PNAS (2020), arXiv:1911.01429] 31

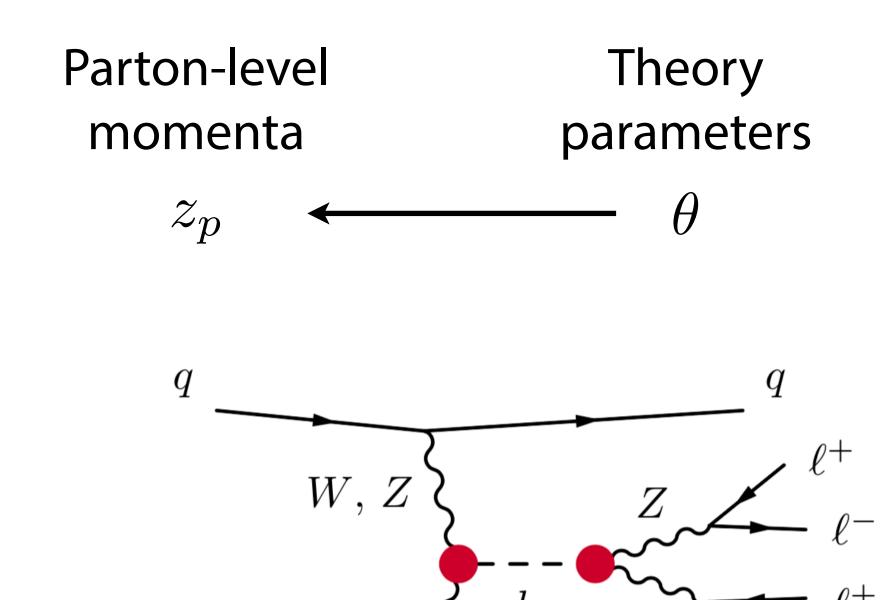
Science is replete with high-fidelity simulators

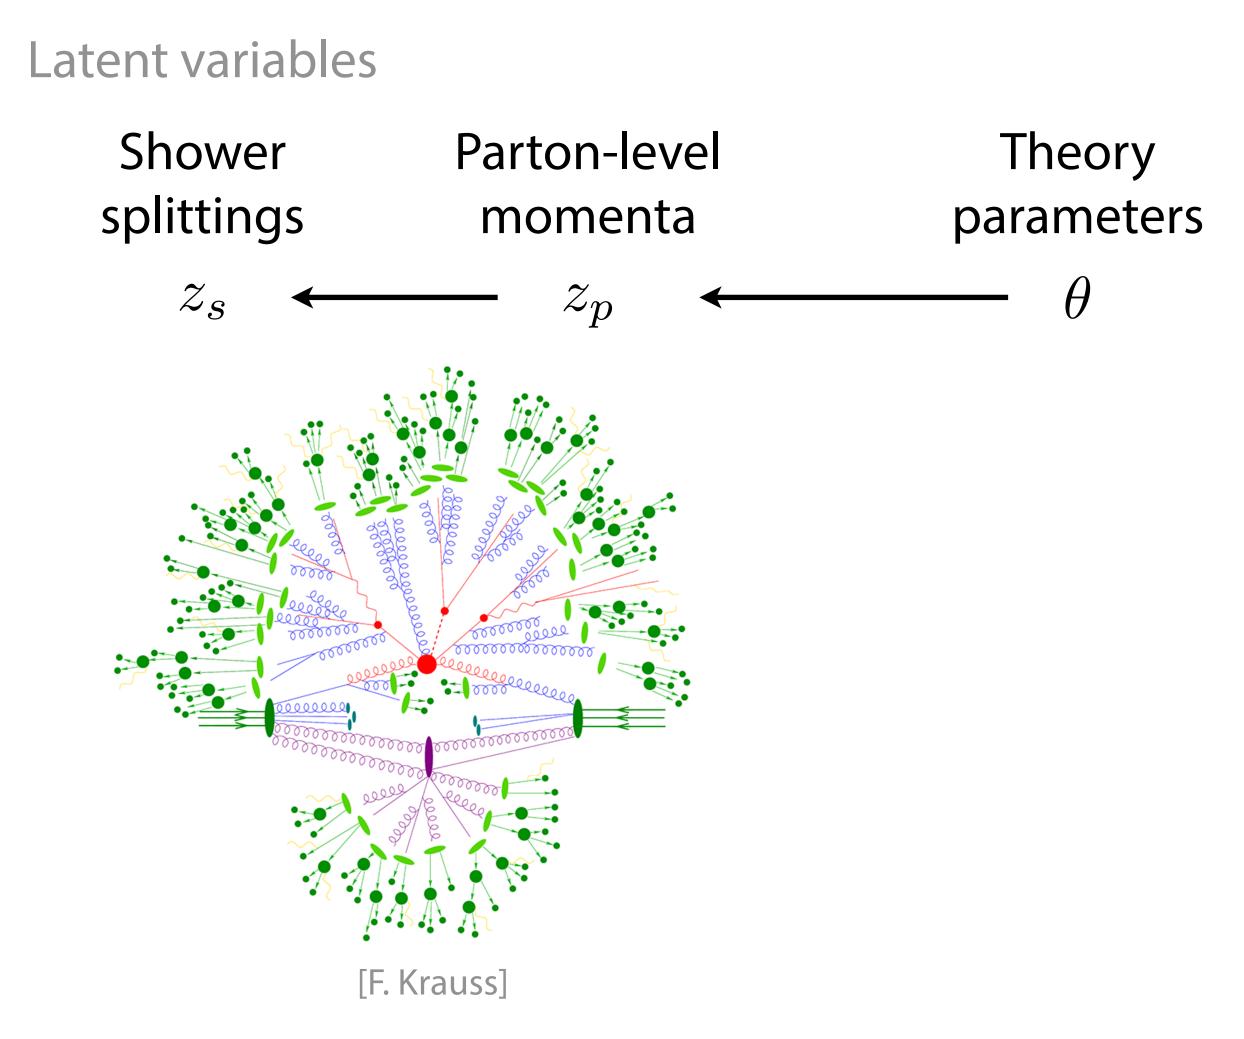


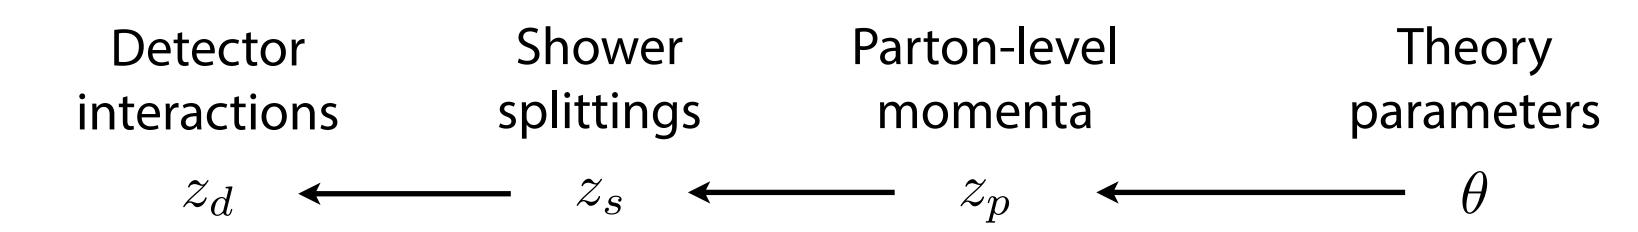
Unfortunately, these simulators are poorly suited for inference.

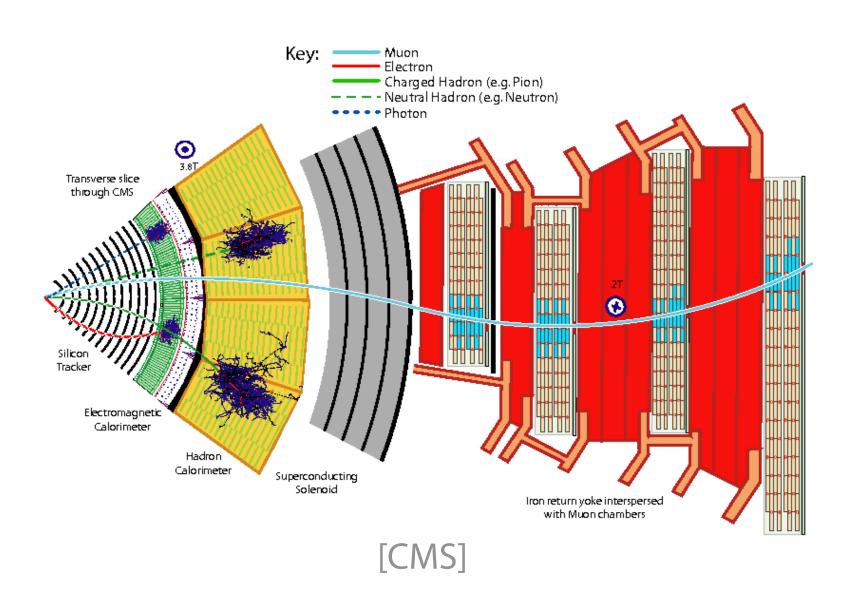
Often involve many latent variables, non-differentiable components, or they are just too slow.

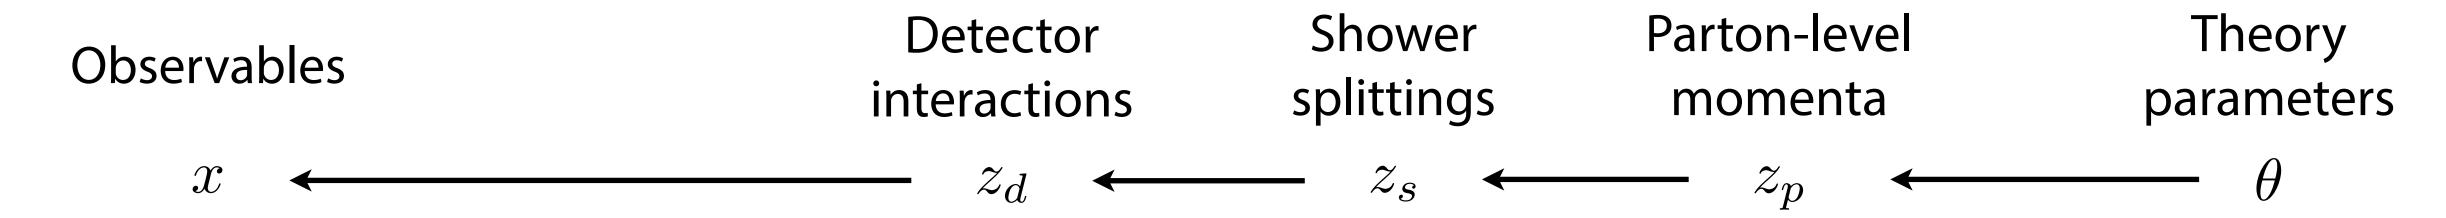
Theory parameters

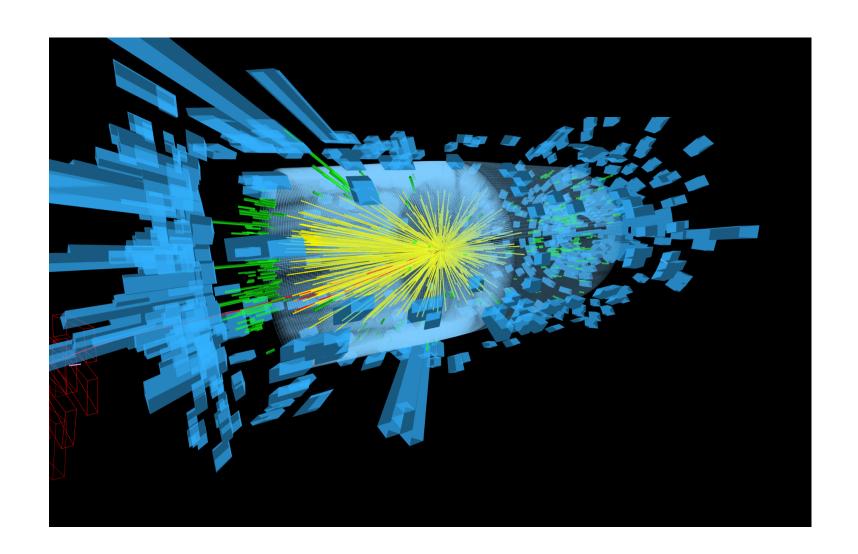


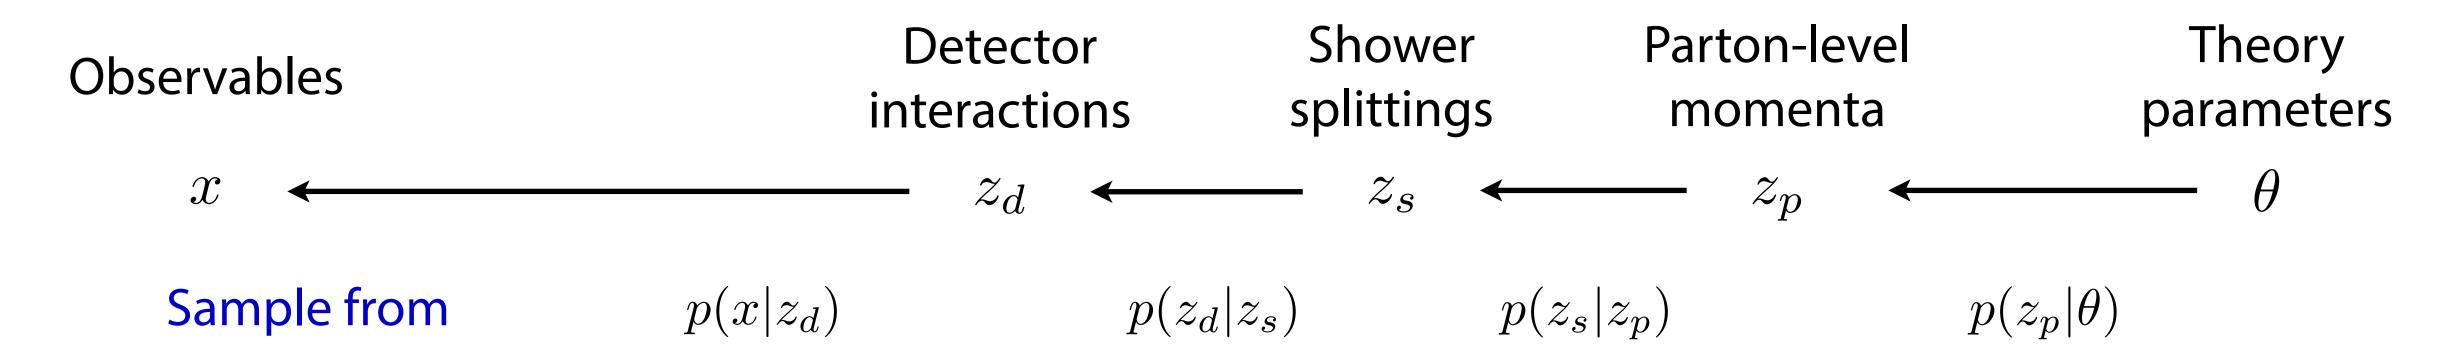






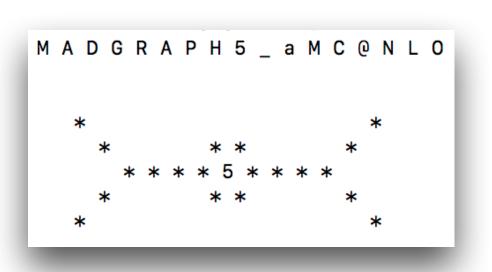


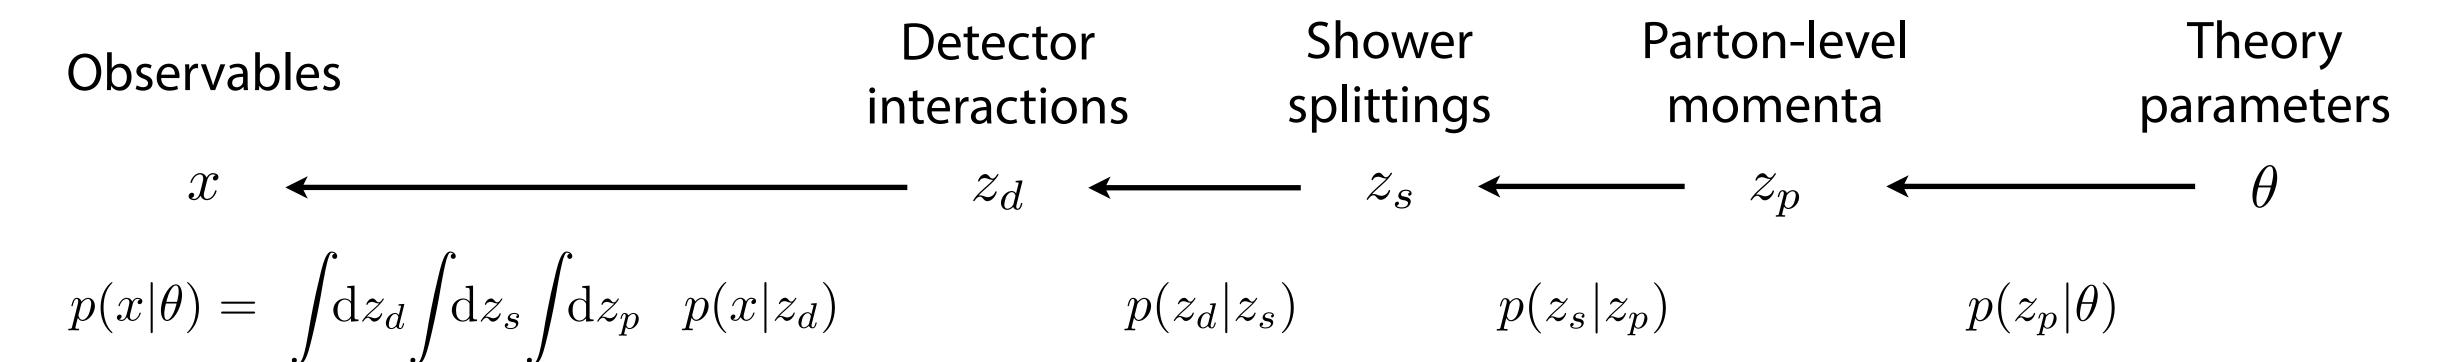




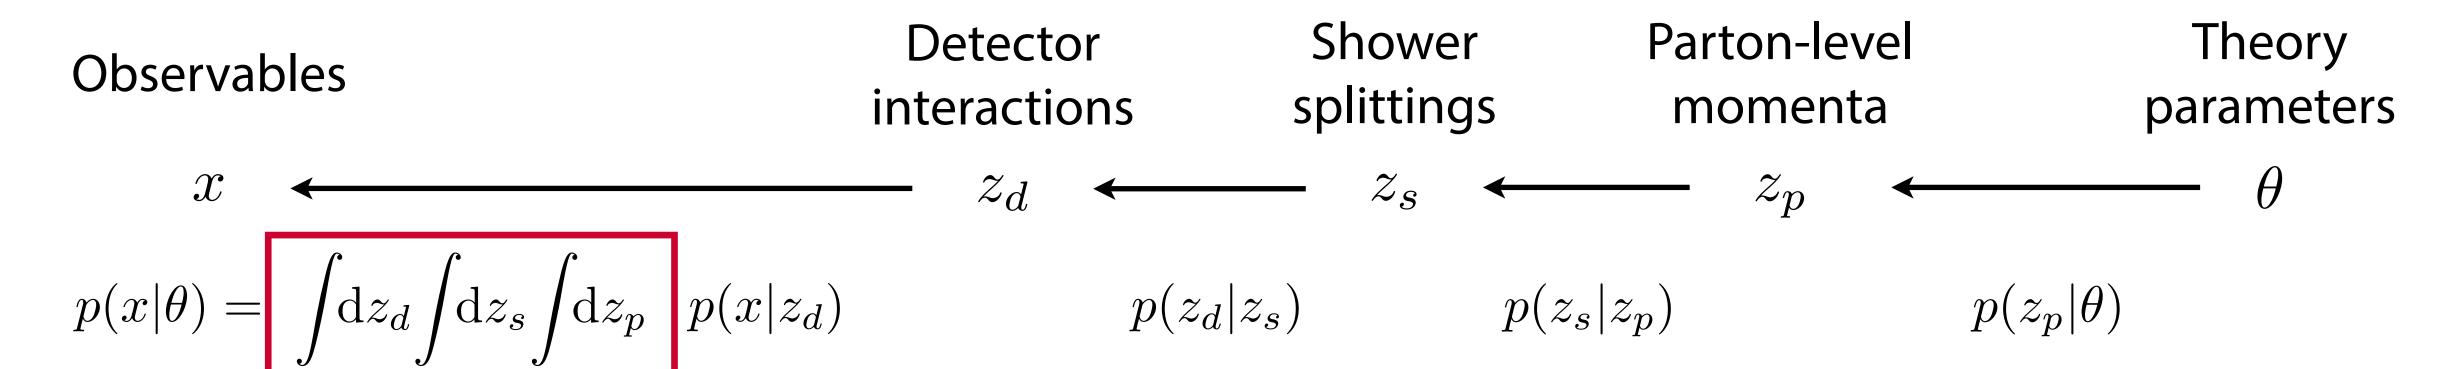








Latent variables

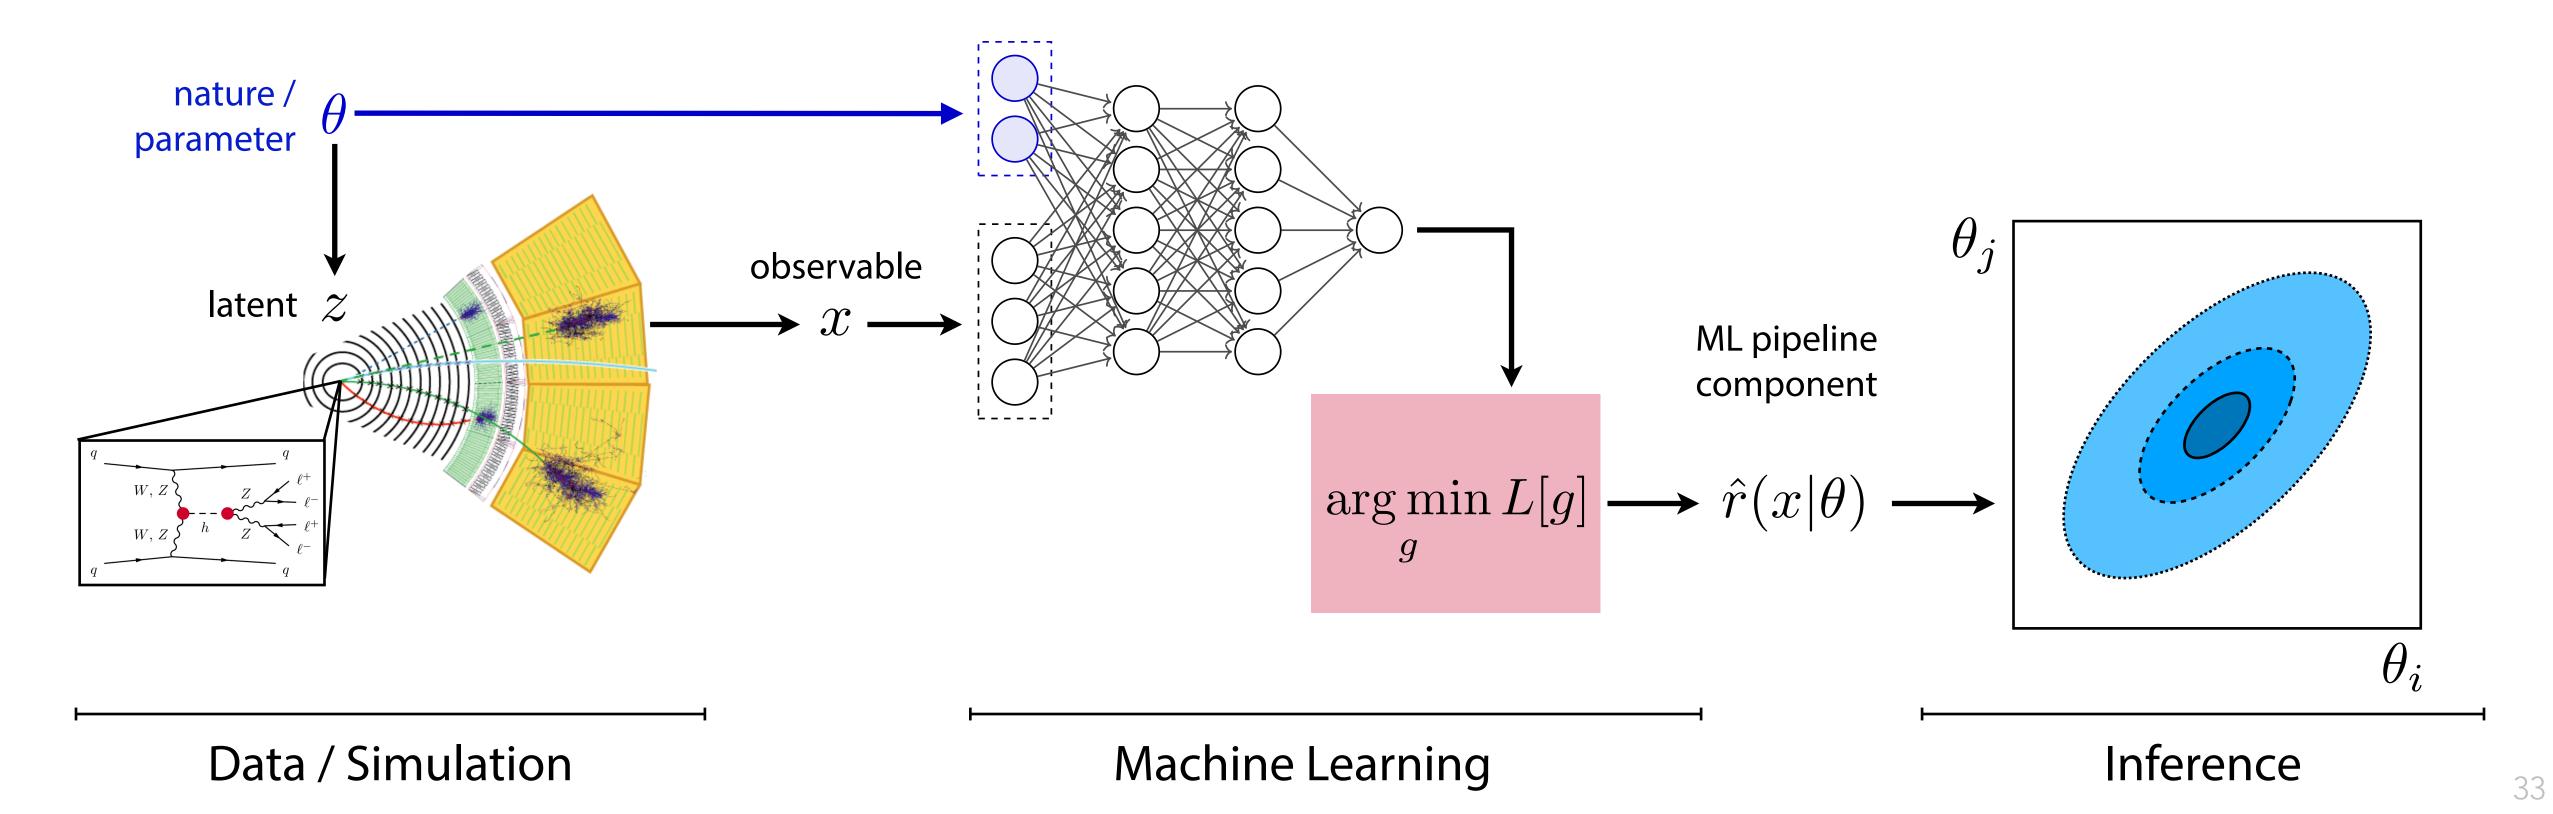


It's infeasible to calculate the integral over this enormous space!

Simulation-Based Inference

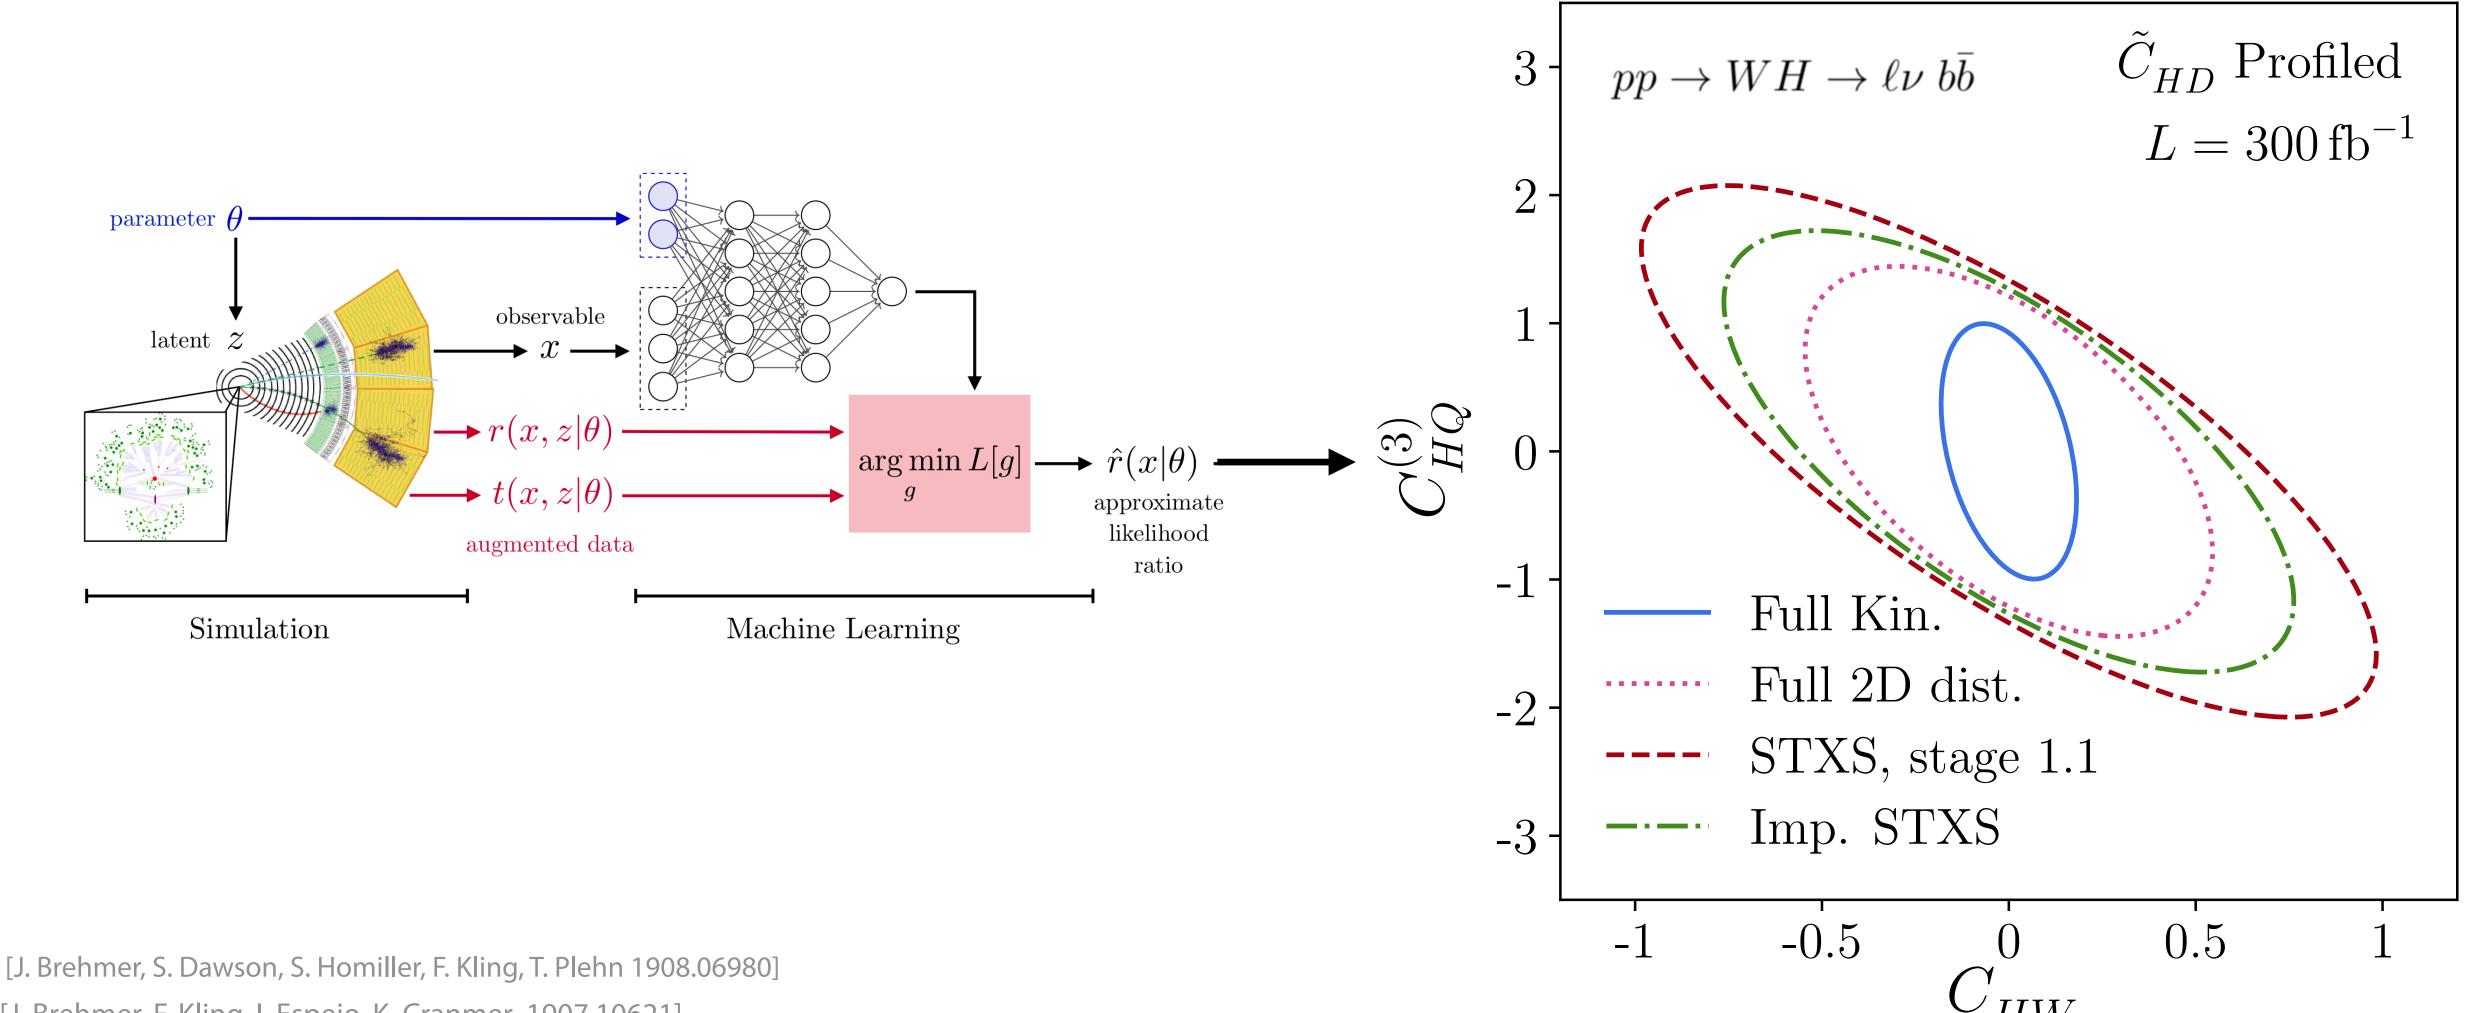
Deep learning and neural density estimation are effective at learning approximate surrogates for the likelihood and posterior, **revolutionizing principled statistical inference in science!**

• Removes the need for hand-engineered summary statistics that sacrifice power



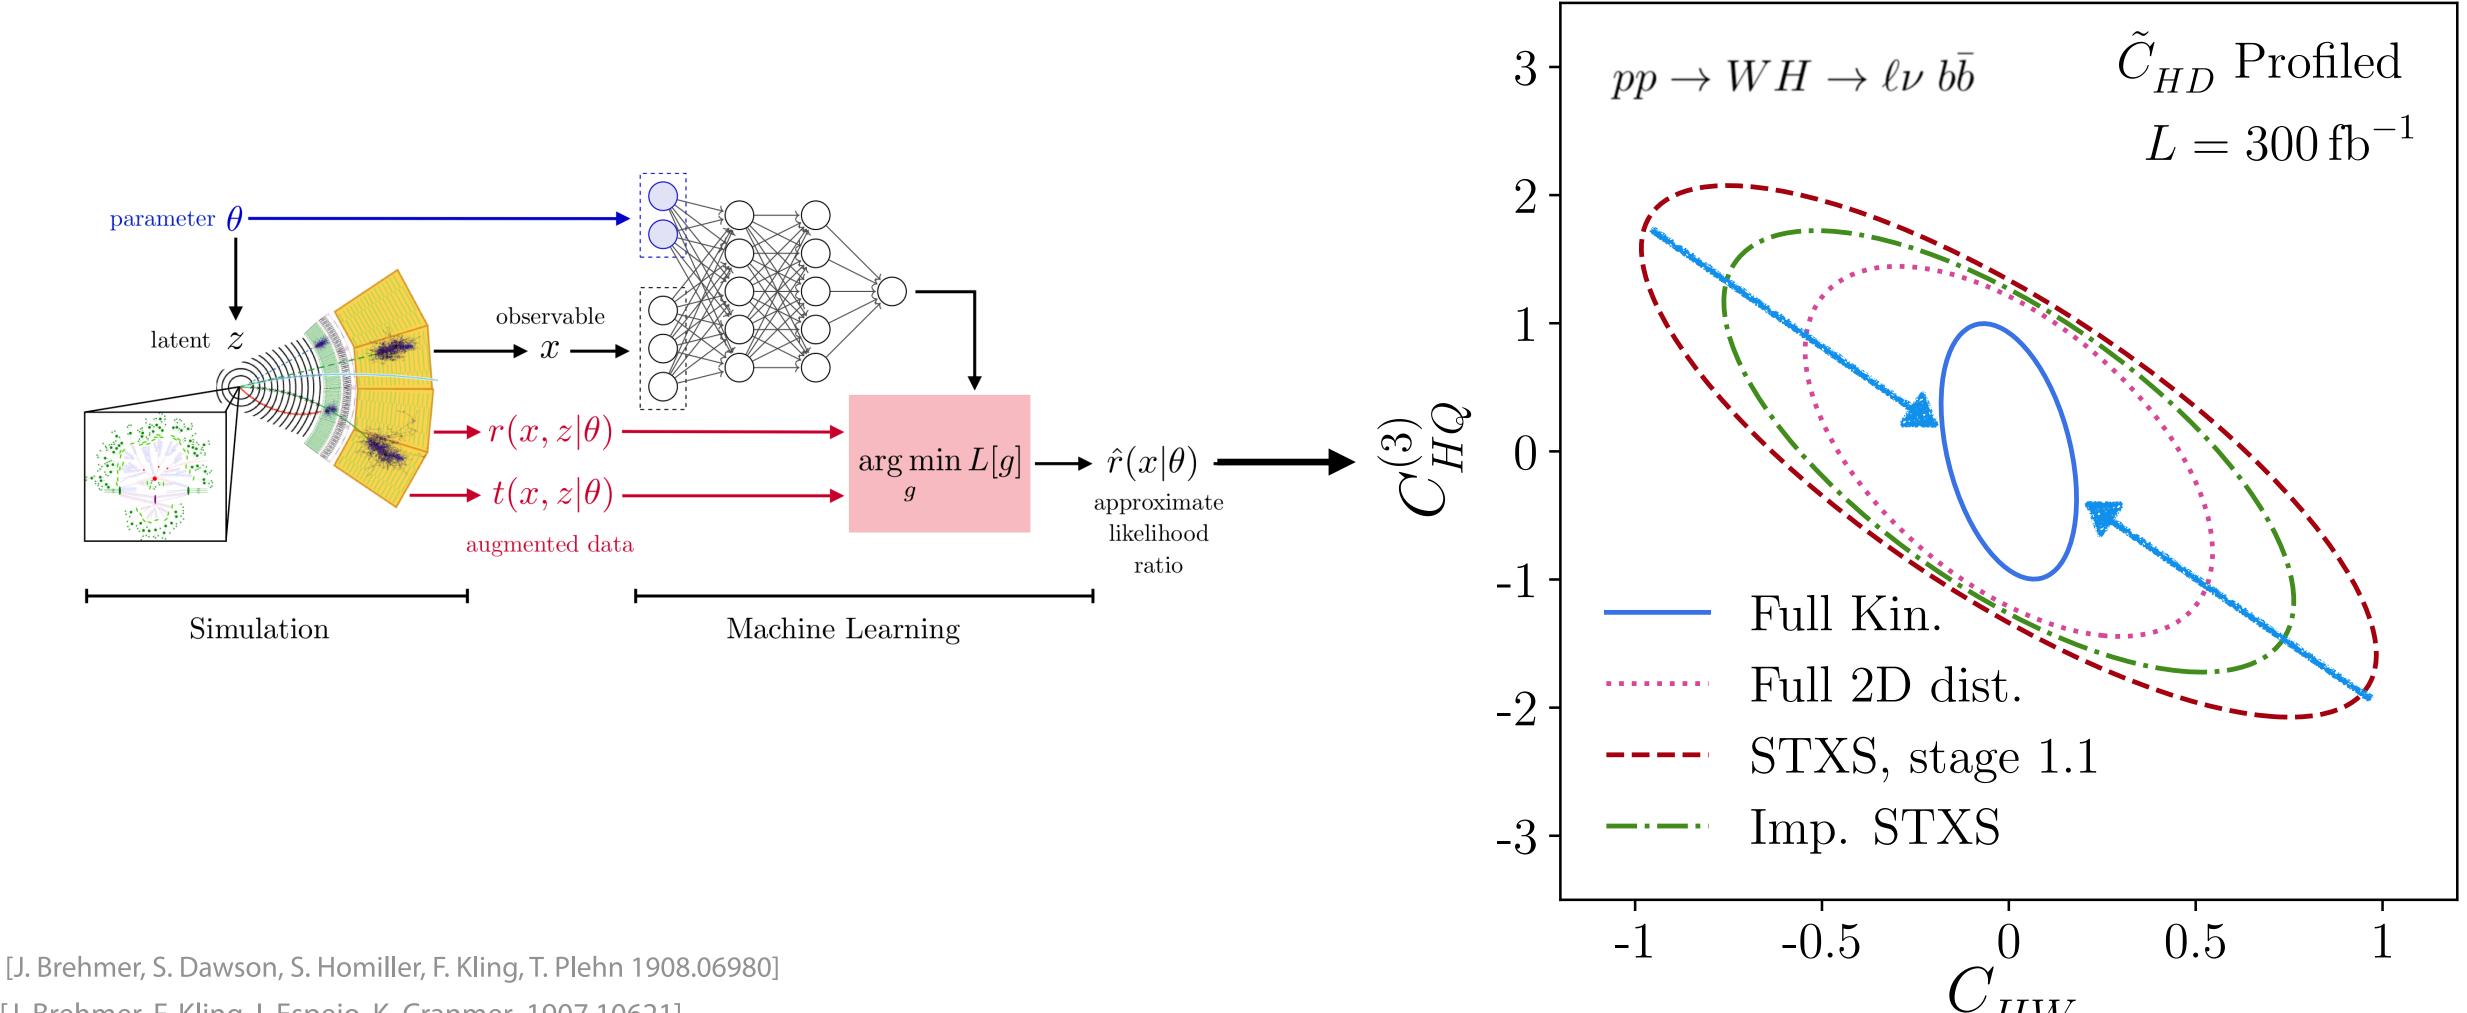
Impact on Studies of The Higgs Boson

Potential for massive gains in precision of a flagship measurement at the LHC! Equivalent increasing data collected by LHC by several factors

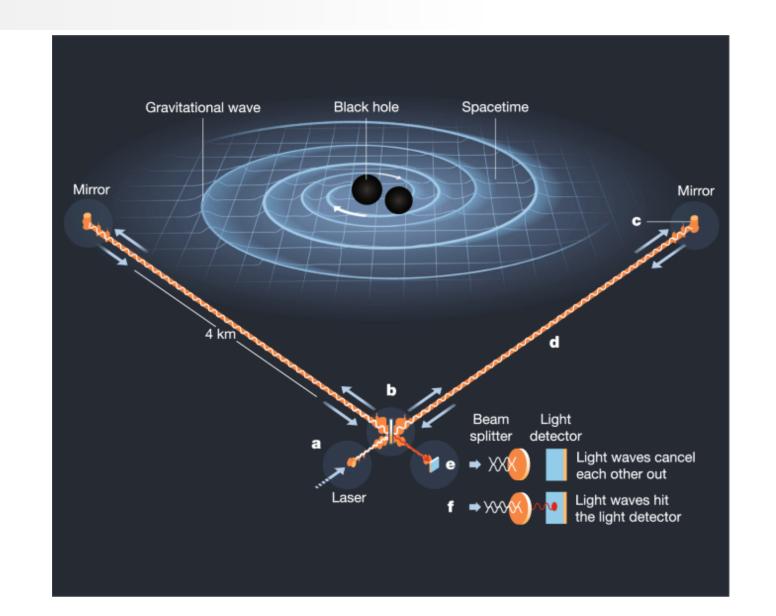


Impact on Studies of The Higgs Boson

Potential for massive gains in precision of a flagship measurement at the LHC! Equivalent increasing data collected by LHC by several factors



Gravitational Wave Astronomy



Real-time gravitational-wave science with neural posterior estimation

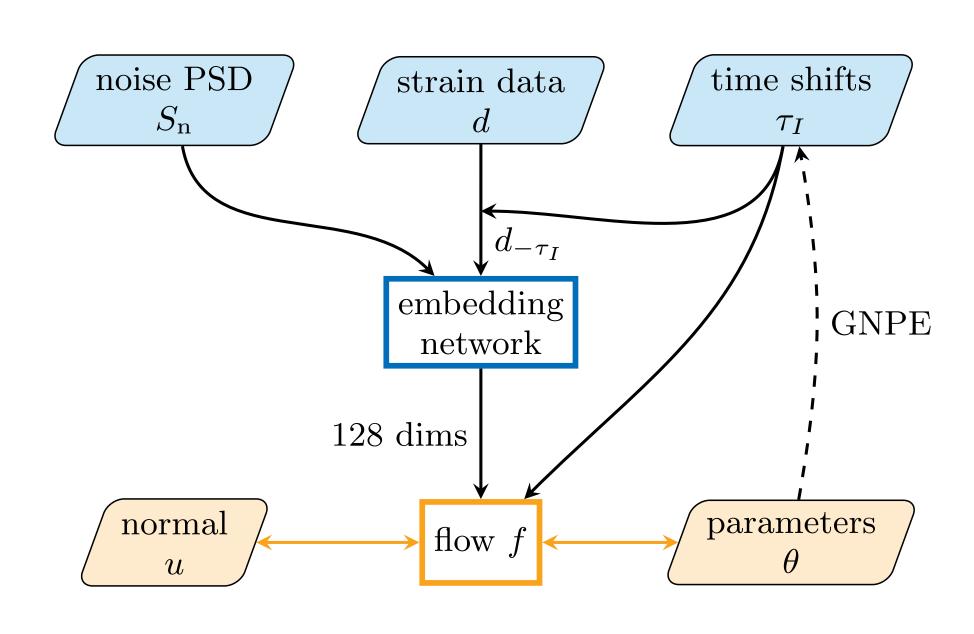
Maximilian Dax,^{1,*} Stephen R. Green,^{2,†} Jonathan Gair,^{2,‡} Jakob H. Macke,^{1,3} Alessandra Buonanno,^{2,4} and Bernhard Schölkopf¹

¹Max Planck Institute for Intelligent Systems, Max-Planck-Ring 4, 72076 Tübingen, Germany

²Max Planck Institute for Gravitational Physics (Albert Einstein Institute), Am Mühlenberg 1, 14476 Potsdam, Germany

³Machine Learning in Science, University of Tübingen, 72076 Tübingen, Germany

⁴Department of Physics, University of Maryland, College Park, MD 20742, USA



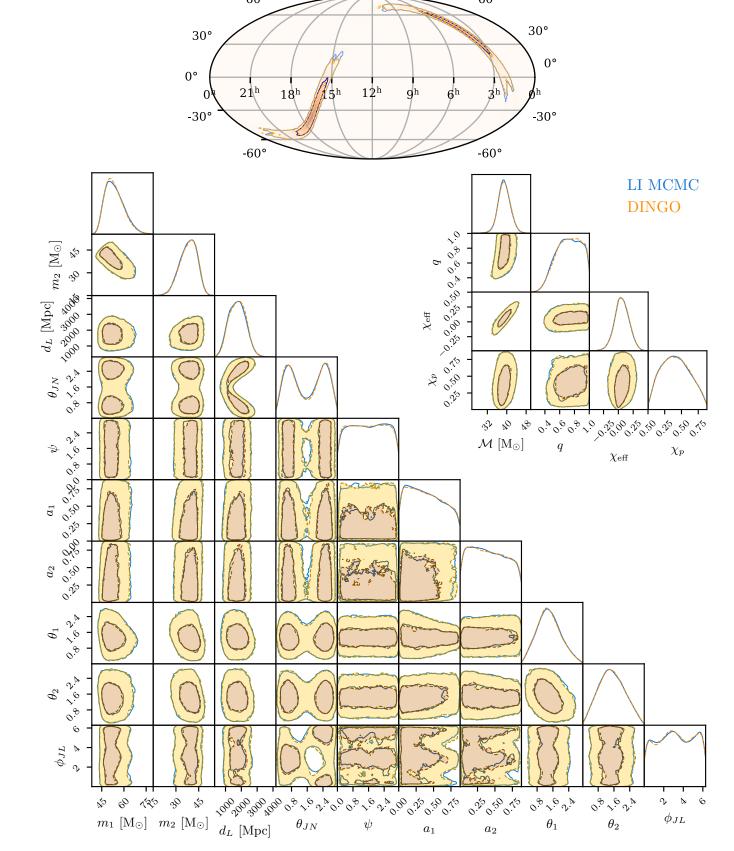


Figure 15. GW170823.

SBI at EuCAIFCon 2024

Simulation-based Inference was well represented at EuCAIFCon!





Biophysical Journal

Supports open access

ABSTRACT VOLUME 122, ISSUE 3, SUPPLEMENT 1, 140A, FEBRUARY 10, 2023

Simulation-based inference of single-molecule force spectroscopy

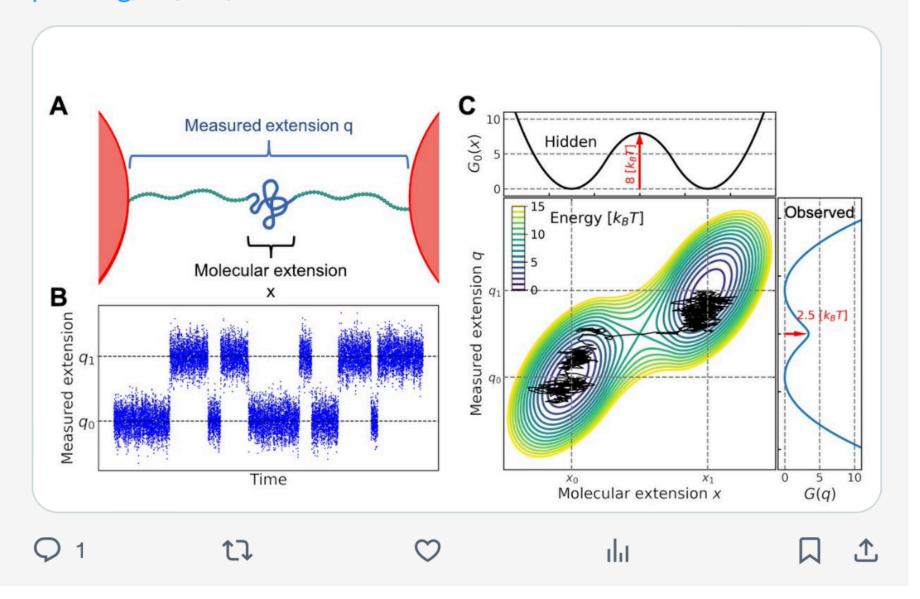
Roberto Covino • Pilar Cossio • Lars Dingeldein



Roberto Covino @CovinoLab · Sep 22, 2022

Because to measure the molecule's motion we have to couple it to a device that is much larger and slower than the molecules itself. The device reports only indirectly on the molecule. Disentangling the device from the molecule is hard.

pnas.org/doi/abs/10.107...





Roberto Covino @CovinoLab · Sep 22, 2022

Using standard statistical techniques is made difficult by the fact that even minimal ideal models lead to mathematically intractable likelihoods that are computationally expensive.

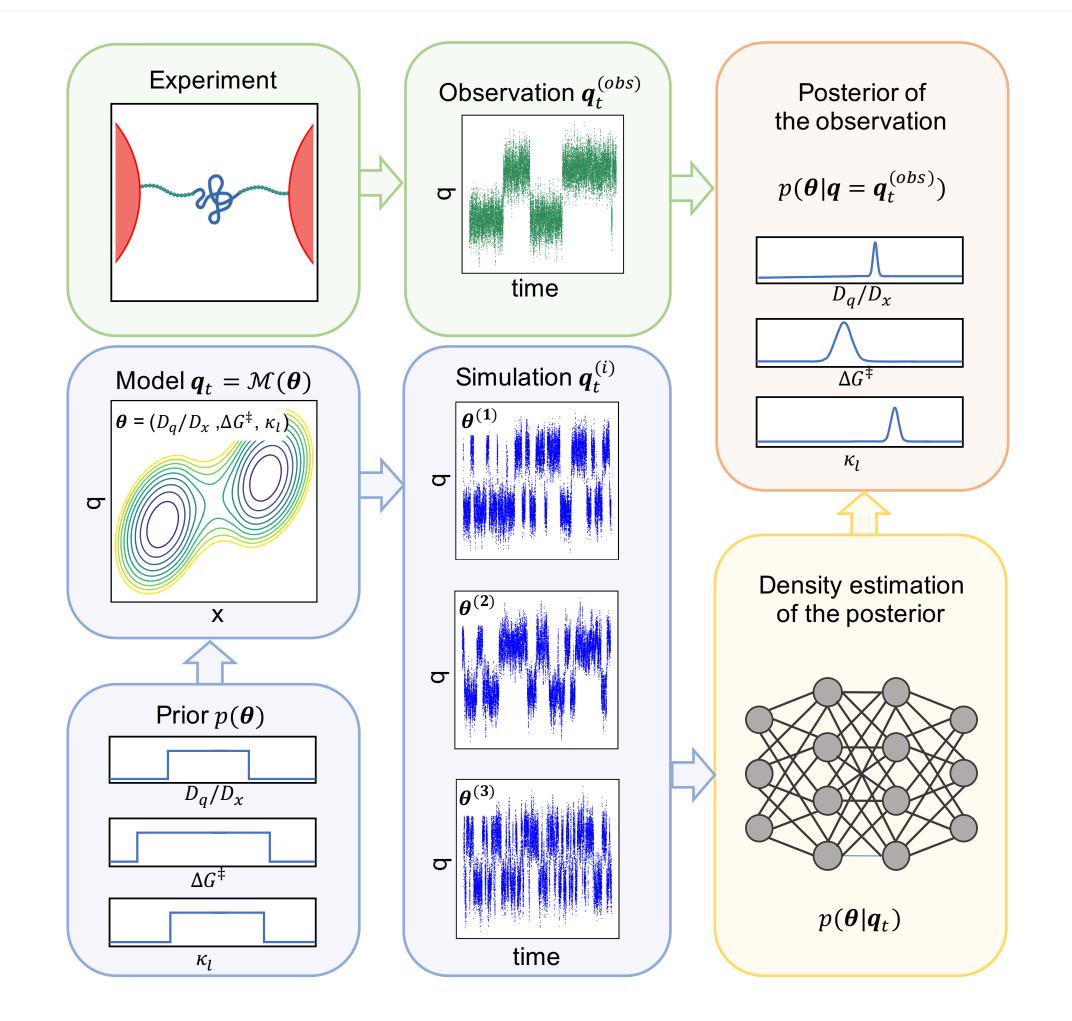


Fig. 2. Workflow of Simulation-Based Inference of smFS experiments. Given an experimental observation $q_t^{(\text{obs})}$, we want to explain it with a parametric mechanistic model $\mathcal{M}(\theta)$ encoded in a simulator. The parameters θ are drawn from the prior distribution $p(\theta)$. The simulator takes parameters drawn from the prior and generates synthetic observations $q_t^{(i)}$. The synthetic observations and corresponding parameters are used to train a conditional density estimator to approximate the posterior $p(\theta|q_t)$. Evaluating the posterior, $p(\theta|q_t = q_t^{(\text{obs})})$, we obtain the most plausible parameters explaining a given observation.

R 7

simulation-based-inference.org

Simulation-based Inference Home Papers Software Jobs Misc About

Papers

Paper

The list is automatically compiled each day. Should you observe any inaccuracies or concerns, kindly bring them to our attention. Additionally, if you believe a new paper aligns with the topic, feel free to submit it.

Visualize the annual growth in the number of publications.

Sort by Category

Total (744) Statistics (195)

Computer Science (102)

Astrophysics (71)

Mathematics (54)

Education (47)

Economics (46)

Physics (33)

Quantitative Biology (31)

Neuroscience (27)

Quantitative Finance (21)

Astronomy (14)

Genetics (13)

Epidemiology (11)

Engineering (10)

Medicine (8)

Geography (8)

Social Science (7)

Evolutionary biology (6)

Ecology (5)

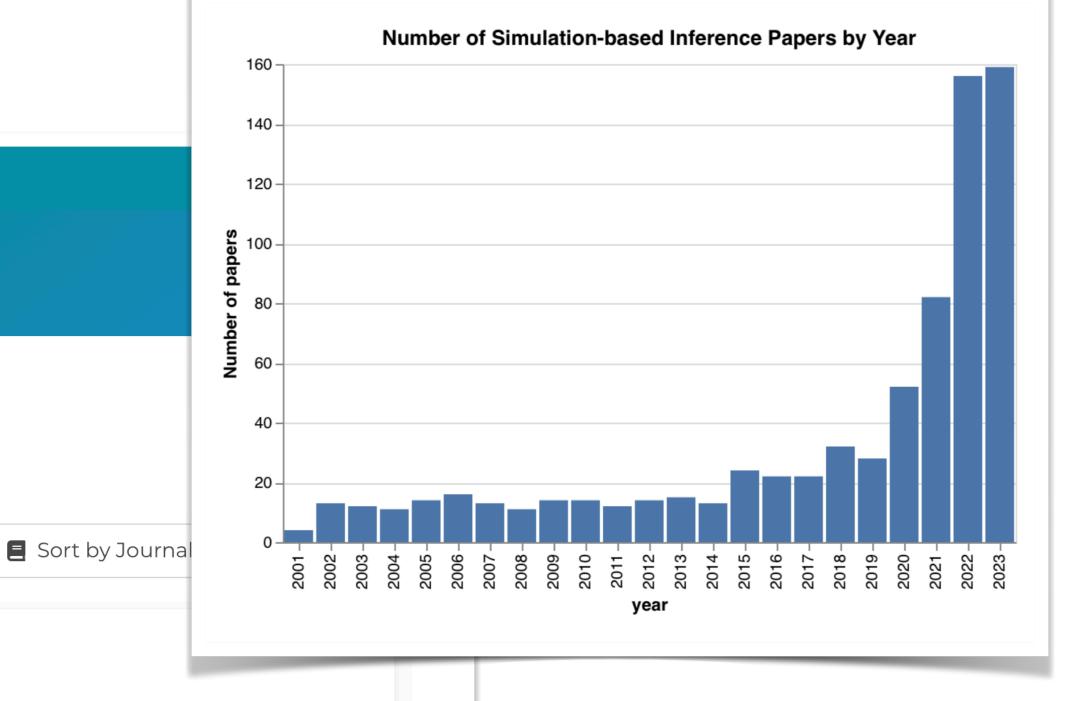
Cognitive Science (4)

Robotics (4)

Systems biology (4)

Statistics

- Simulation based stacking, Y Yao, BRS Blancard, J Domke arXiv preprint arXiv:2310.17009, 2023 arxiv.org
- Calibrating Neural Simulation-Based Inference with Differentiable Coverage Probability, M Falkiewicz, N Takeishi, I Shekhzadeh... arXiv preprint arXiv ..., 2023 arxiv.org
- Simulation-based Inference with the Generalized Kullback-Leibler Divergence, BK Miller, M Federici, C Weniger, P Forré arXiv preprint arXiv ..., 2023 arxiv.org
- Simulation-based Inference for Cardiovascular Models, A Wehenkel, J Behrmann, AC Miller, G Sapiro... arXiv preprint arXiv ..., 2023 arxiv.org
- Hierarchical Neural Simulation-Based Inference Over Event Ensembles, L Heinrich, S Mishra-Sharma, C Pollard... arXiv preprint arXiv ..., 2023 arxiv.org
- L-C2ST Local Diagnostics for Posterior Approximations in Simulation-Based Inference, J Linhart, A Gramfort, PLC Rodrigues arXiv preprint arXiv:2306.03580, 2023 arxiv.org
- Learning Robust Statistics for Simulation-based Inference under Model Misspecification, D Huang, A Bharti, A Souza, L Acerbi... arXiv preprint arXiv ..., 2023 arxiv.org
- Generalized Bayesian Inference for Scientific Simulators via Amortized Cost Estimation, R Gao, M Deistler, JH Macke arXiv preprint arXiv:2305.15208, 2023 arxiv.org
- Variational Inference with Coverage Guarantees, Y Patel, D McNamara, J Loper, J Regier... arXiv preprint arXiv ..., 2023
 arxiv.org
- Generalised likelihood profiles for models with intractable likelihoods, DJ Warne, OJ Maclaren, EJ Carr, MJ Simpson...
 arXiv preprint arXiv ..., 2023 arxiv.org
- Neural Likelihood Surfaces for Spatial Processes with Computationally Intensive or Intractable Likelihoods, J Walchessen, A Lenzi, M Kuusela arXiv preprint arXiv:2305.04634, 2023 arxiv.org
- Ralancing Simulation-based Inference for Conservative Posteriors A Delaunov RK Miller P Forré C Weniger arXiv



Computational Microscope

The strong force: Quantum Chromodynamics (QCD)

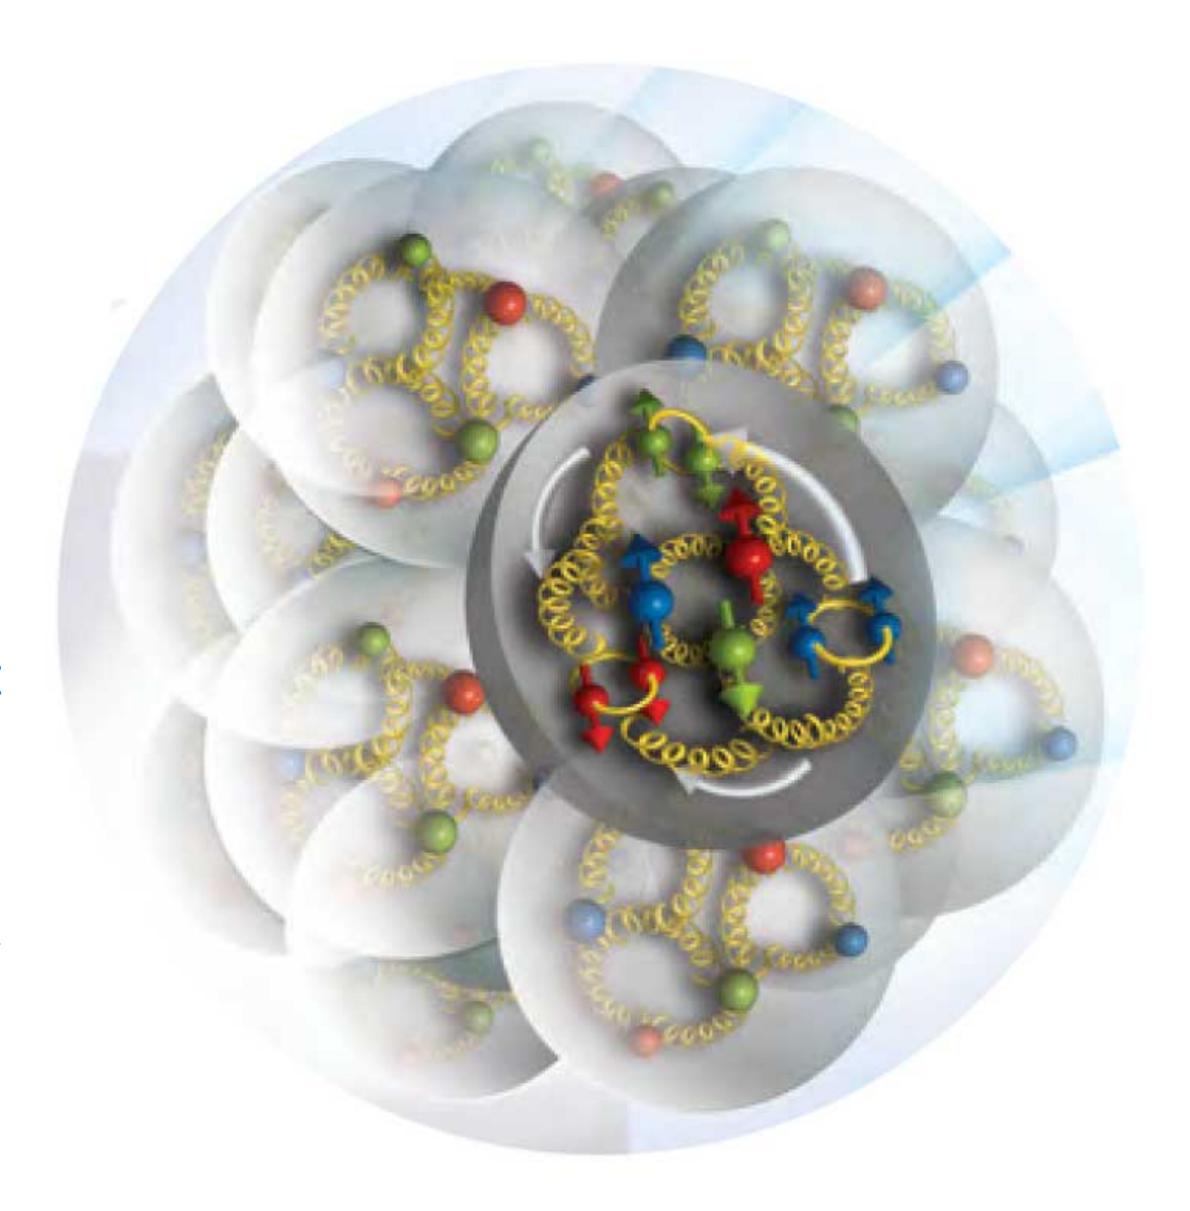
The strong nuclear force is one of the four fundamental forces.

It is described by Quantum Chromodynamics (QCD)

QCD describes how quarks and gluons interact

Emergent phenomena:

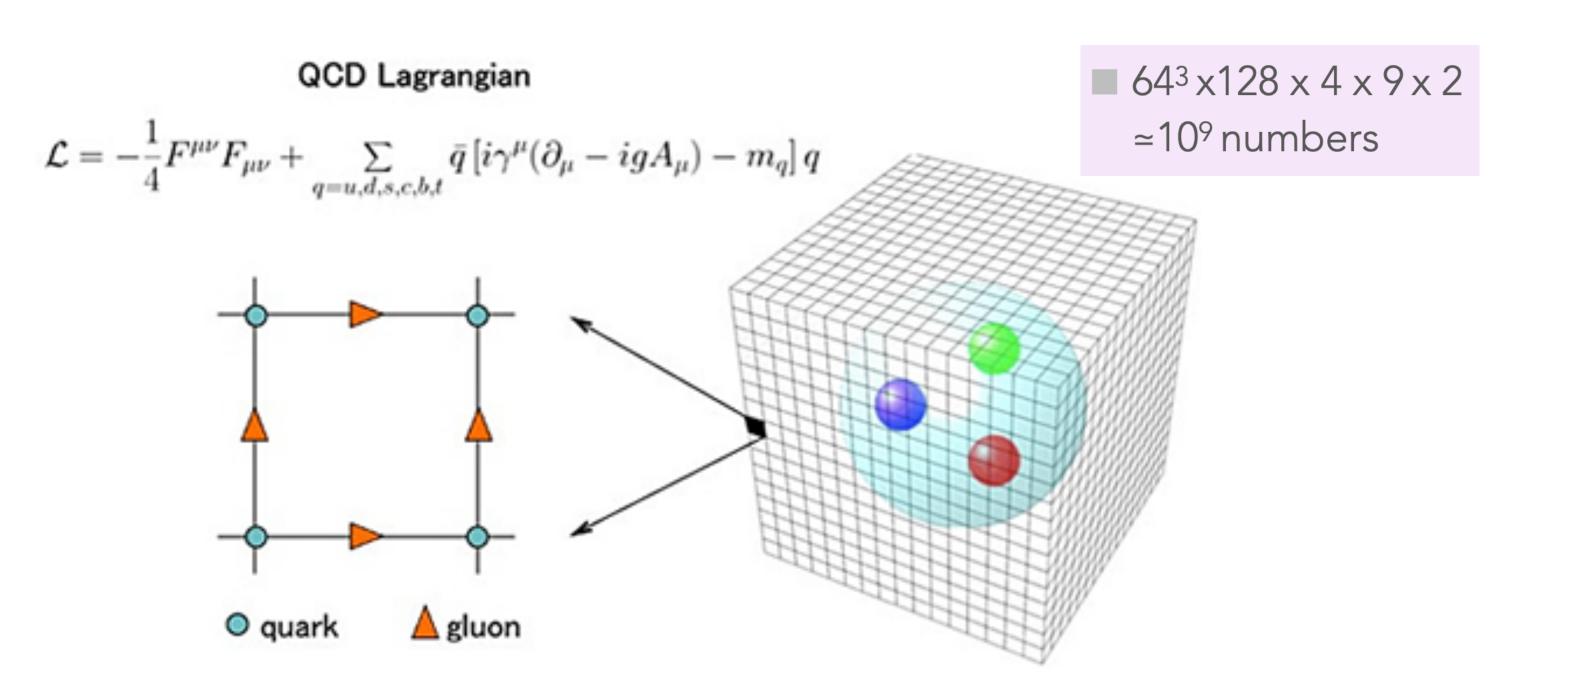
Quarks and gluons form protons, neutrons, etc.



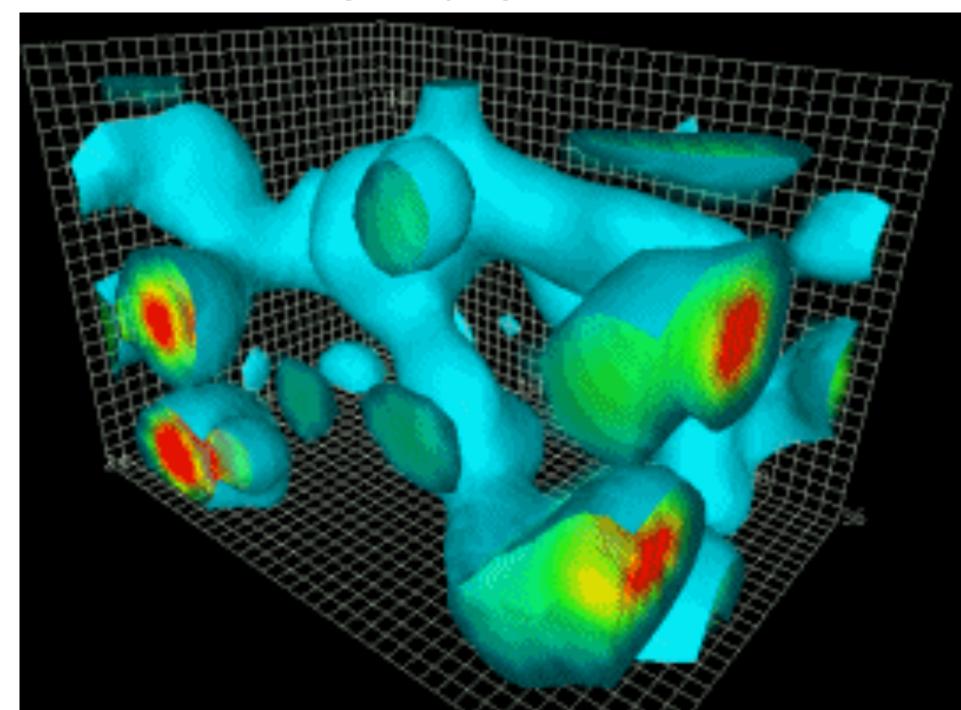
Lattice Field Theory

Lattice field theory is a computational approach to studying interacting field theory on a discretized space-time lattice.

Each link on the lattice has data corresponding to the symmetry group of the theory. For the strong force (QCD) each link has a 3x3 unitary matrix.



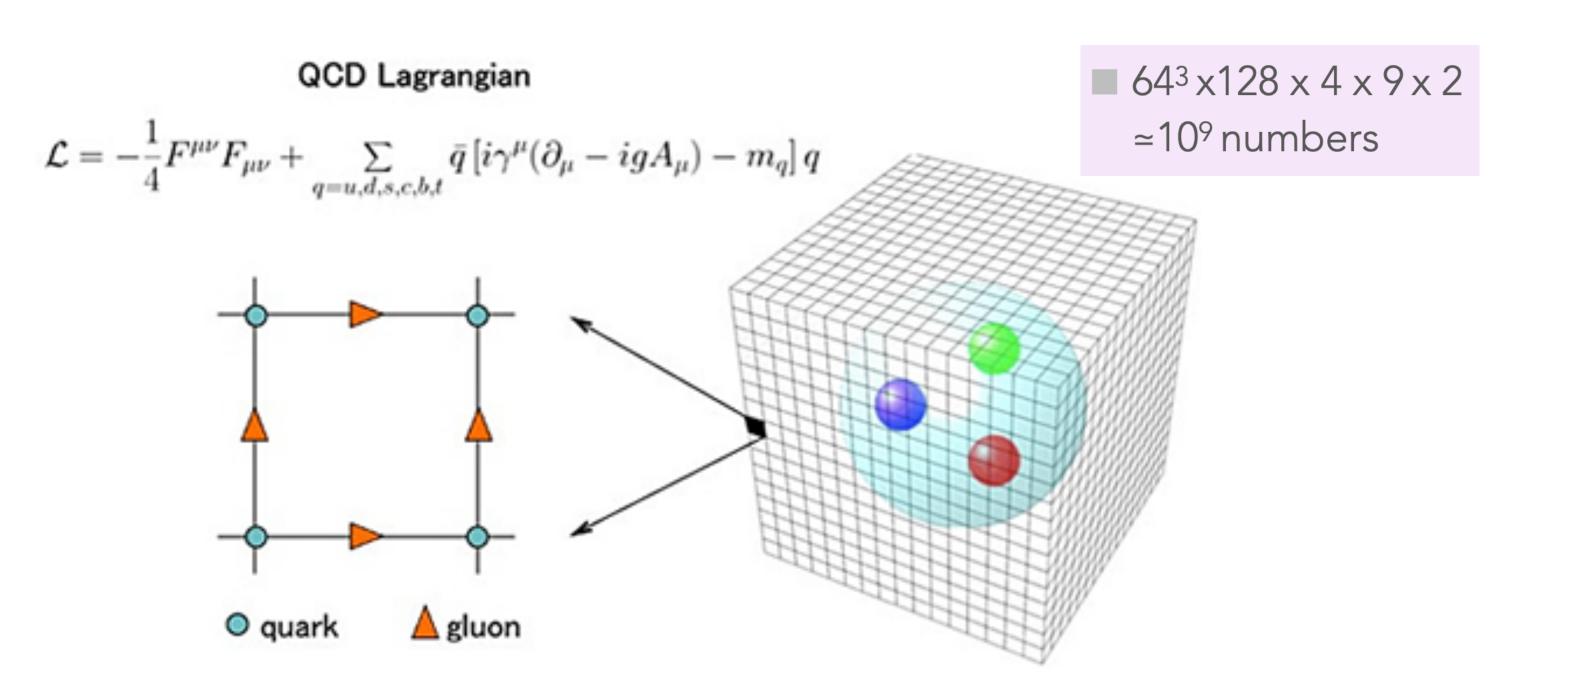
This animation is a single configuration of the lattice. Think of a 4-d image playing like a movie.



Lattice Field Theory

Lattice field theory is a computational approach to studying interacting field theory on a discretized space-time lattice.

Each link on the lattice has data corresponding to the symmetry group of the theory. For the strong force (QCD) each link has a 3x3 unitary matrix.



This animation is a single configuration of the lattice. Think of a 4-d image playing like a movie.

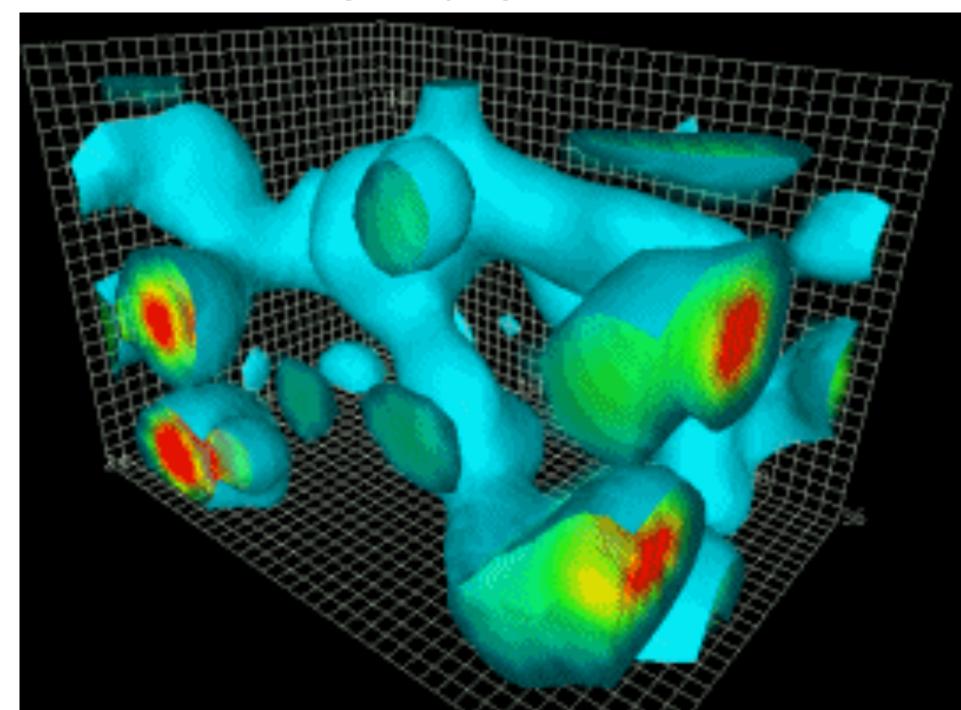
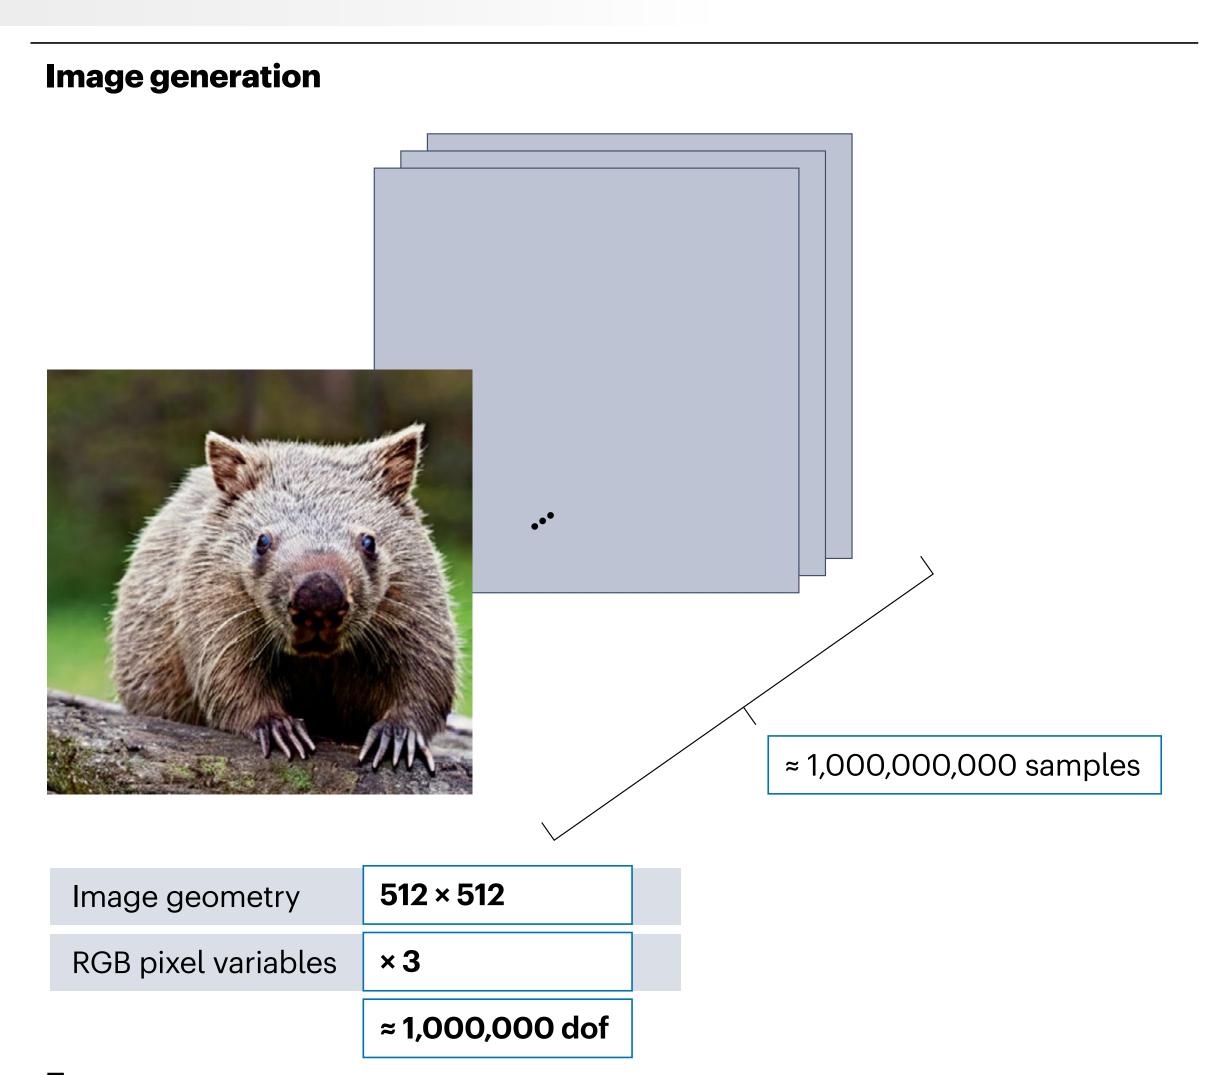


Image vs. Lattice Quantum Fields



Target

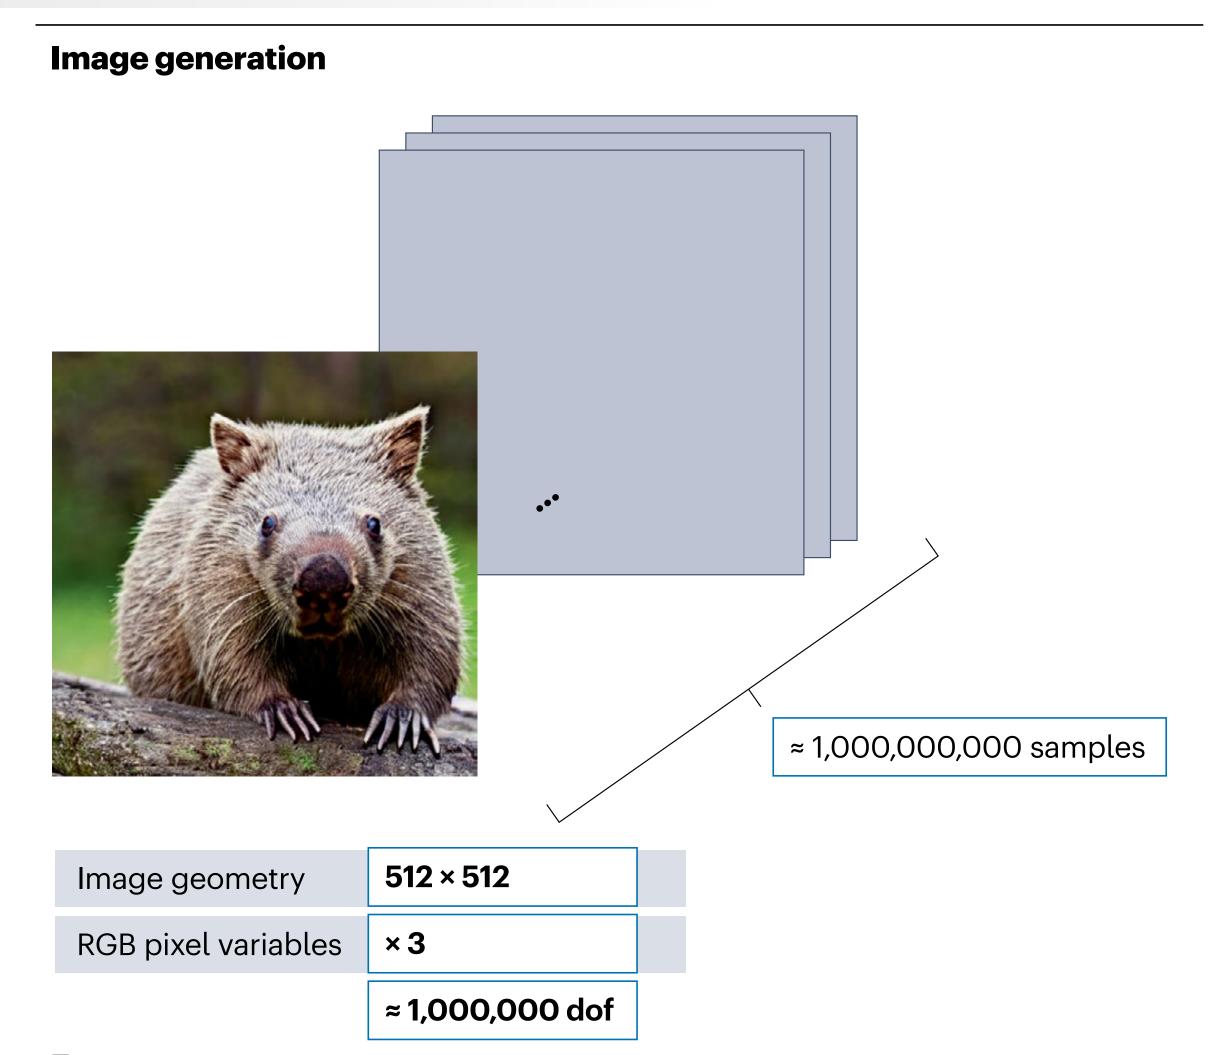
Subjective high quality per sample

Symmetries

Few approximate symmetries (for example, reflection, small translations)

Image vs. Lattice Quantum Fields

sampling motivated by lattice quatum chromodynamics

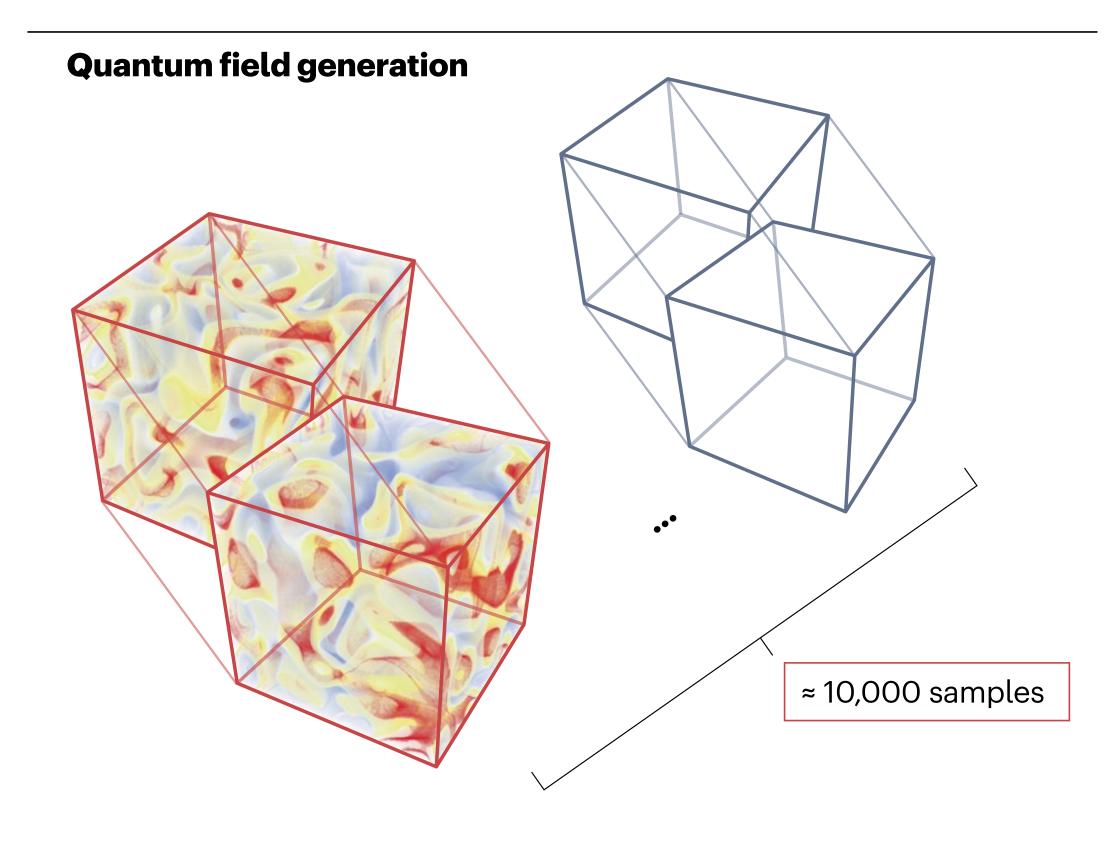


Target

Subjective high quality per sample

Symmetries

Few approximate symmetries (for example, reflection, small translations)



Lattice geometry	256 × 256 × 256 × 512
SU(3) link variables	×4×8
	≈ 100,000,000 dof

Target

Objective distribution $p(U) = e^{-S(U)}/Z$

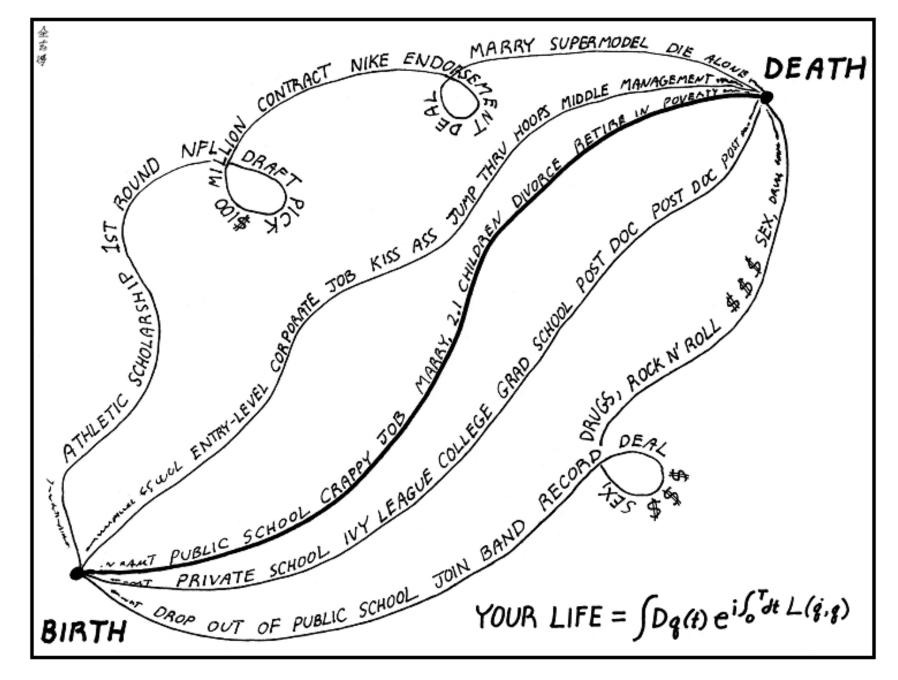
Symmetries

High-dimensional exact symmetries (for example, translations, gauge symmetry)

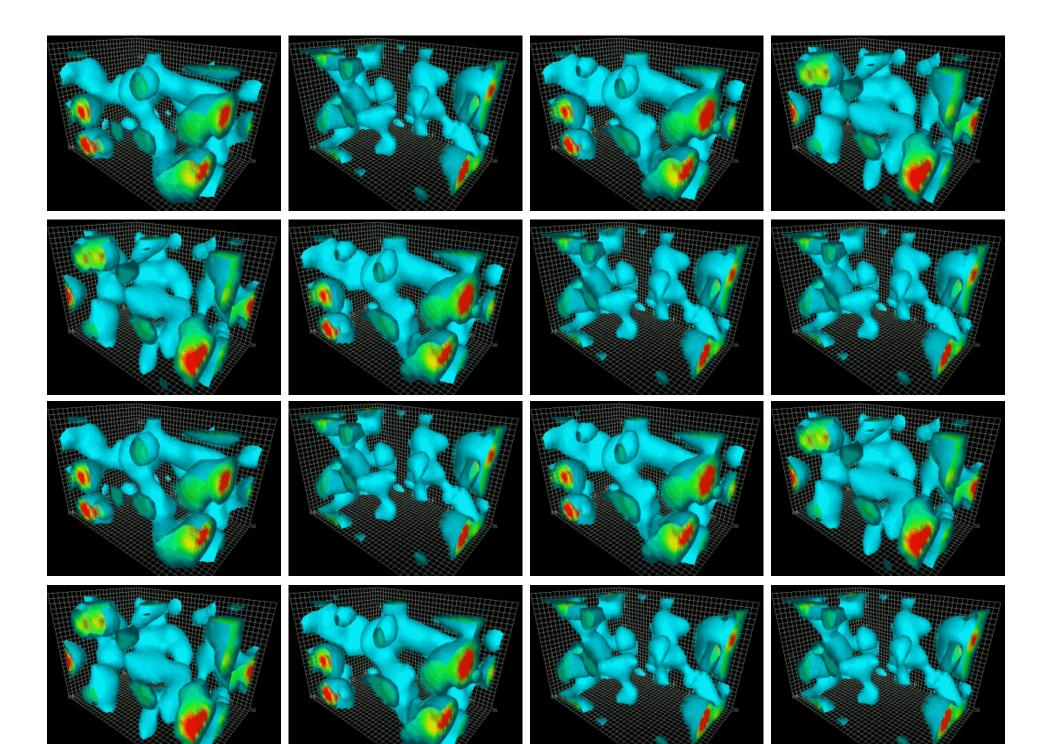
Distribution over configurations

We don't want just a single "image" (lattice configuration), we want to **sample** the high-dimensional distribution of configurations predicted by the theory.

- Path integral: each "path" is a sample from distribution of lattice configurations path ~exp(-Action[path])
- Predictions are **expectations** of quantum operators w.r.t. this distribution.
- Hamiltonian Monte Carlo was invented for this problem, but it has limitations.



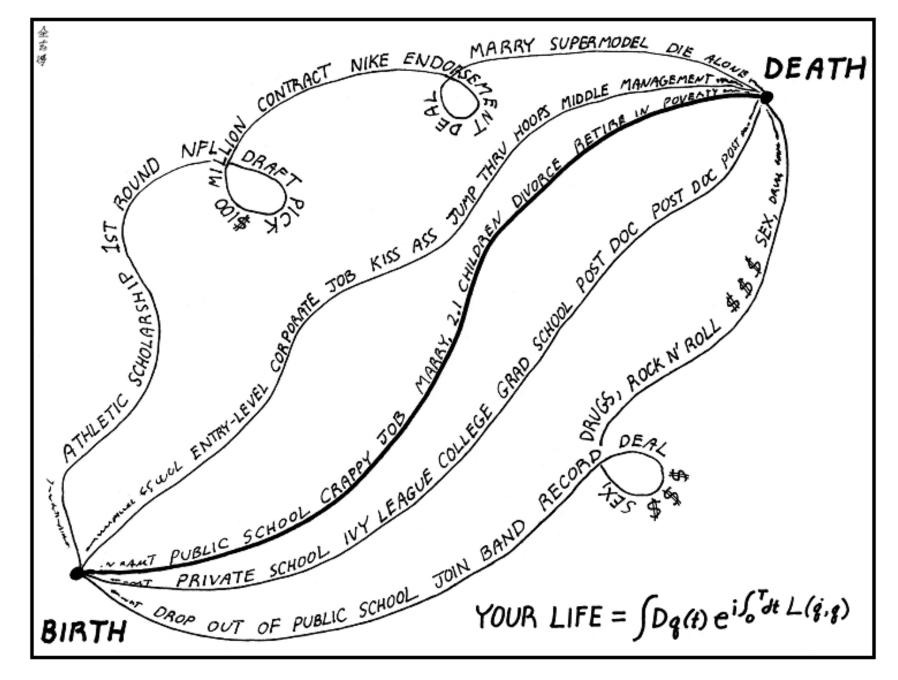
The Path Integral Formulation of Your Life



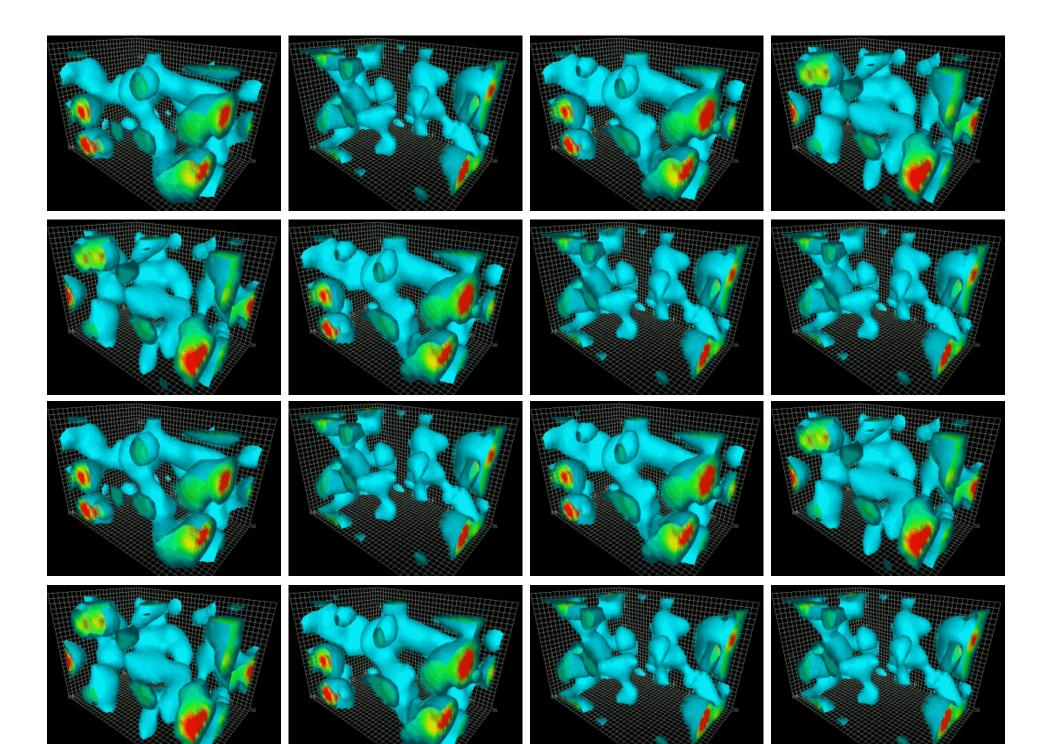
Distribution over configurations

We don't want just a single "image" (lattice configuration), we want to **sample** the high-dimensional distribution of configurations predicted by the theory.

- Path integral: each "path" is a sample from distribution of lattice configurations path ~exp(-Action[path])
- Predictions are **expectations** of quantum operators w.r.t. this distribution.
- Hamiltonian Monte Carlo was invented for this problem, but it has limitations.



The Path Integral Formulation of Your Life



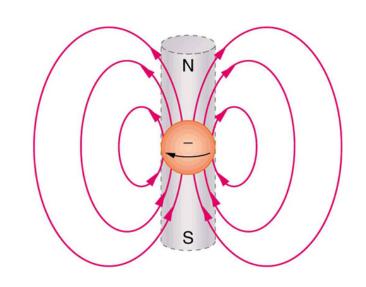
Predictions are taken seriously

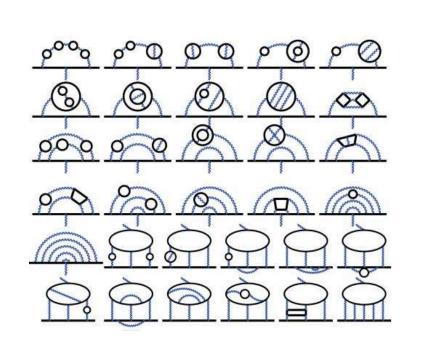
Magnetic moment of the electron: (torque an electron feels in a magnetic field) $a_e = (g-2)/2$

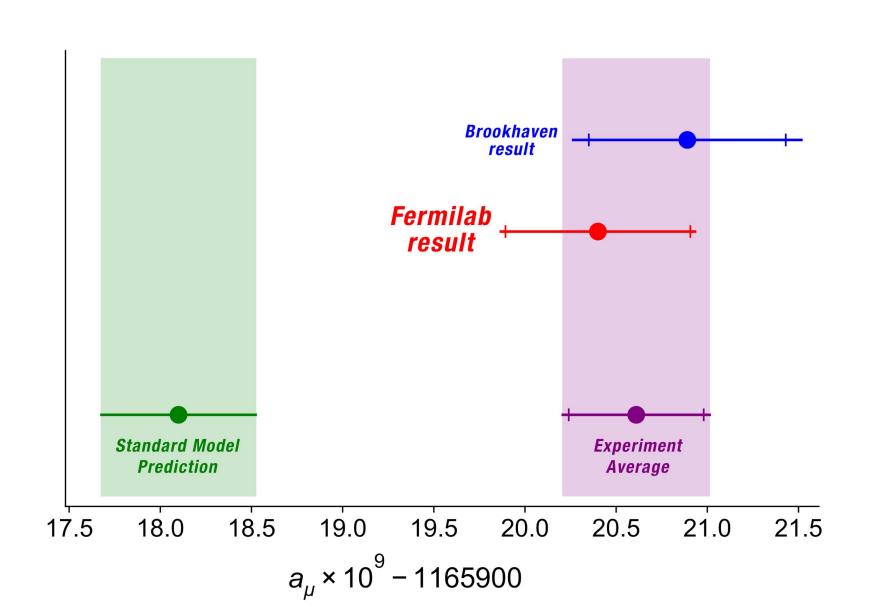


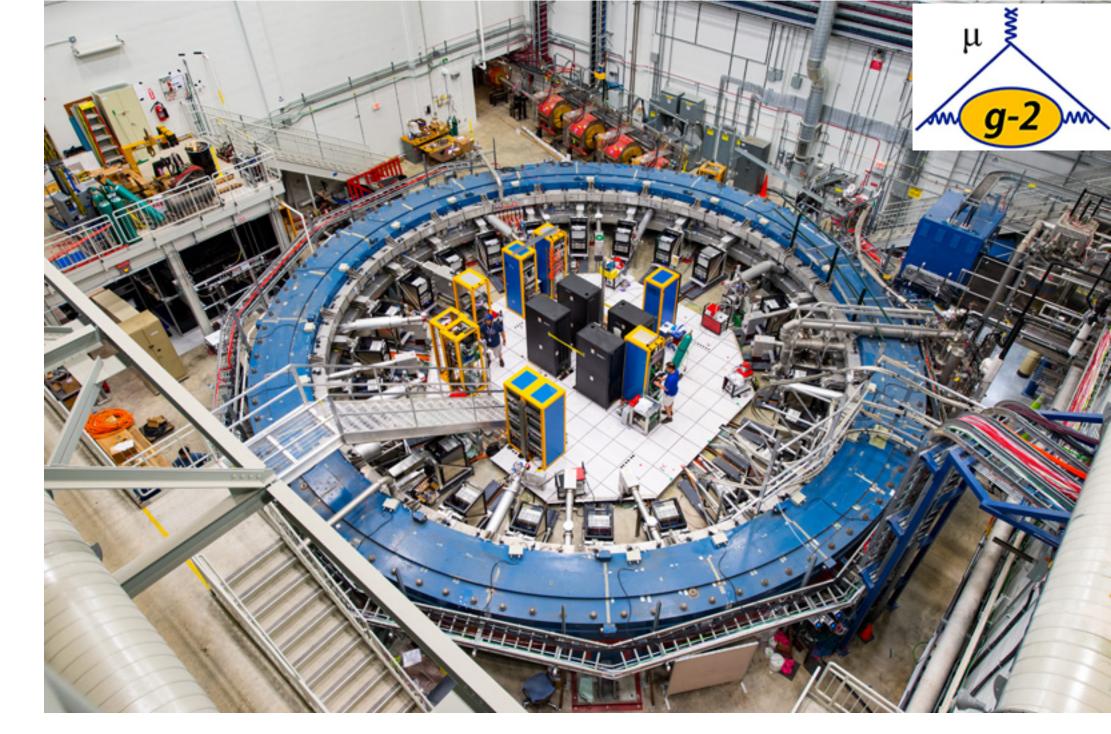
Theory $a_e = 0.001159652181643(764)$

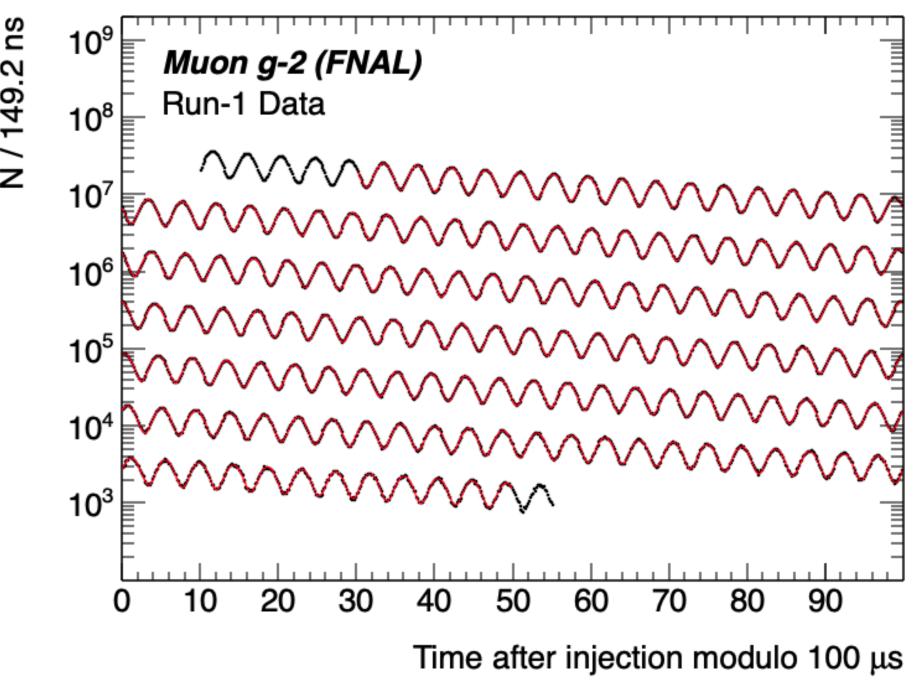
Exp. $a_e = 0.00115965218073(28)$











Flows for LQCD

Basic idea:

- use normalizing flows to approximate the target Boltzmann distribution.
- Train using reverse KL[q||p] (not samples from the target)
- Sample from the flow instead of traditional Hamiltonian MC

Learned model won't be perfect, but you can **correct** via importance sampling or MCMC procedure

Flow-based generative models for Markov chain Monte Carlo in lattice field theory

M. S. Albergo, ^{1,2,3} G. Kanwar, ⁴ and P. E. Shanahan ^{4,1}

¹Perimeter Institute for Theoretical Physics, Waterloo, Ontario N2L 2Y5, Canada ²Cavendish Laboratories, University of Cambridge, Cambridge CB3 0HE, U.K. ³University of Waterloo, Waterloo, Ontario N2L 3G1, Canada ⁴Center for Theoretical Physics, Massachusetts Institute of Technology, Cambridge, MA 02139, U.S.A.

A Markov chain update scheme using a machine-learned flow-based generative model is proposed for Monte Carlo sampling in lattice field theories. The generative model may be optimized (trained) to produce samples from a distribution approximating the desired Boltzmann distribution determined by the lattice action of the theory being studied. Training the model systematically improves autocorrelation times in the Markov chain, even in regions of parameter space where standard Markov chain Monte Carlo algorithms exhibit critical slowing down in producing decorrelated updates. Moreover, the model may be trained without existing samples from the desired distribution. The algorithm is compared with HMC and local Metropolis sampling for ϕ^4 theory in two dimensions.

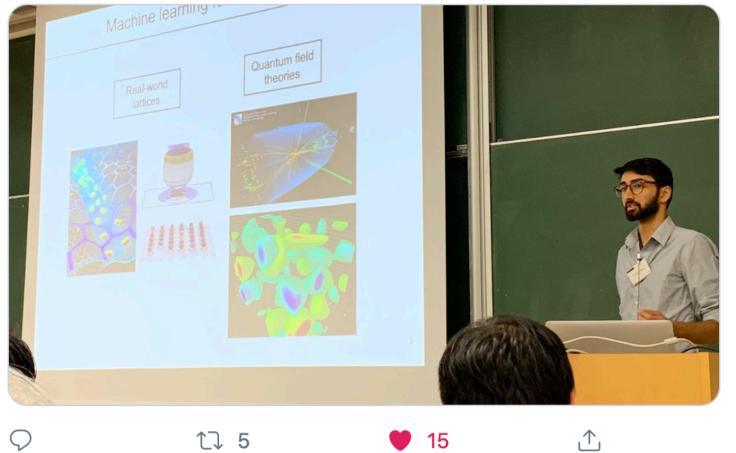


Hire Michael ↑



Enrico Rinaldi @enricesena · Nov 1

Yesterday Gurtej Kanwar told us about machine learning for lattice field theories and exciting progress in Generative Models for gauge theories (collaboration with @DeepMindAI) at #DLAP2019 Today is the last day of this great conference!



Flows for molecular dynamics





RESEARCH

Noé et al., Science **365**, 1001 (2019)

6 September 2019

RESEARCH ARTICLE SUMMARY

MACHINE LEARNING

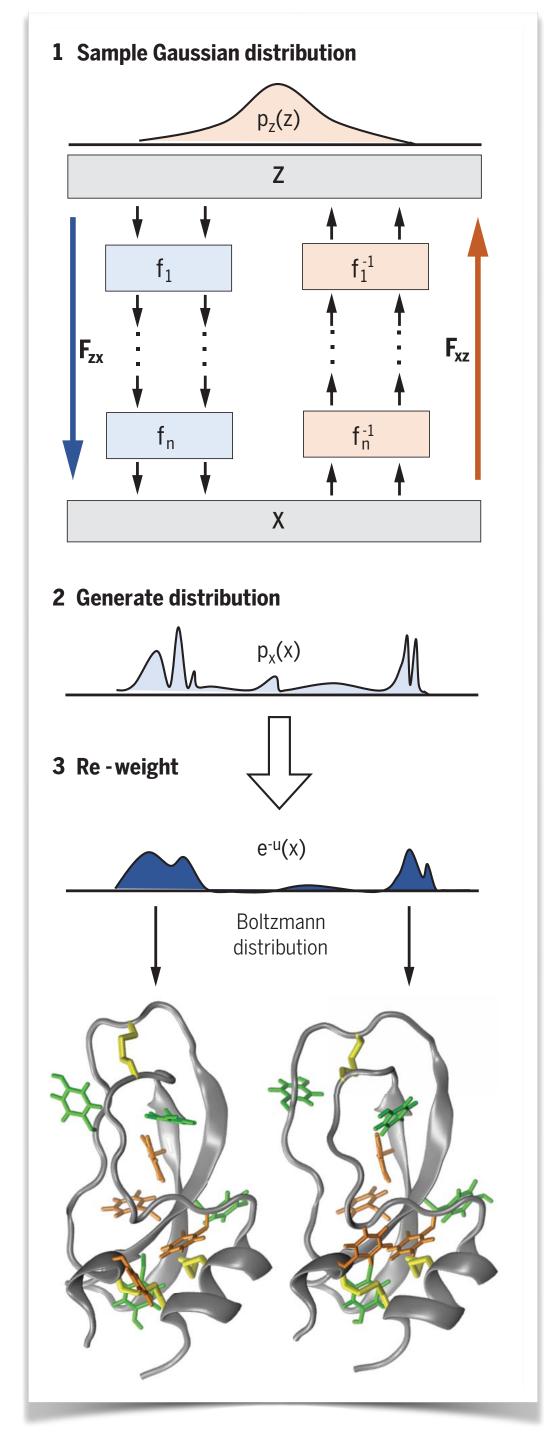
Boltzmann generators: Sampling equilibrium states of many-body systems with deep learning

Frank Noé*†, Simon Olsson*, Jonas Köhler*, Hao Wu

The main approach is thus to start with one configuration, e.g., the folded protein state, and make tiny changes to it over time, e.g., by using Markov-chain Monte Carlo or molecular dynamics (MD). However, these simulations get trapped in metastable (long-lived) states: For example, sampling a single folding or unfolding event with atomistic MD may take a year on a supercomputer.

Same core idea as Boltzmann generators

Boltzmann generators overcome sampling problems between long-lived states.

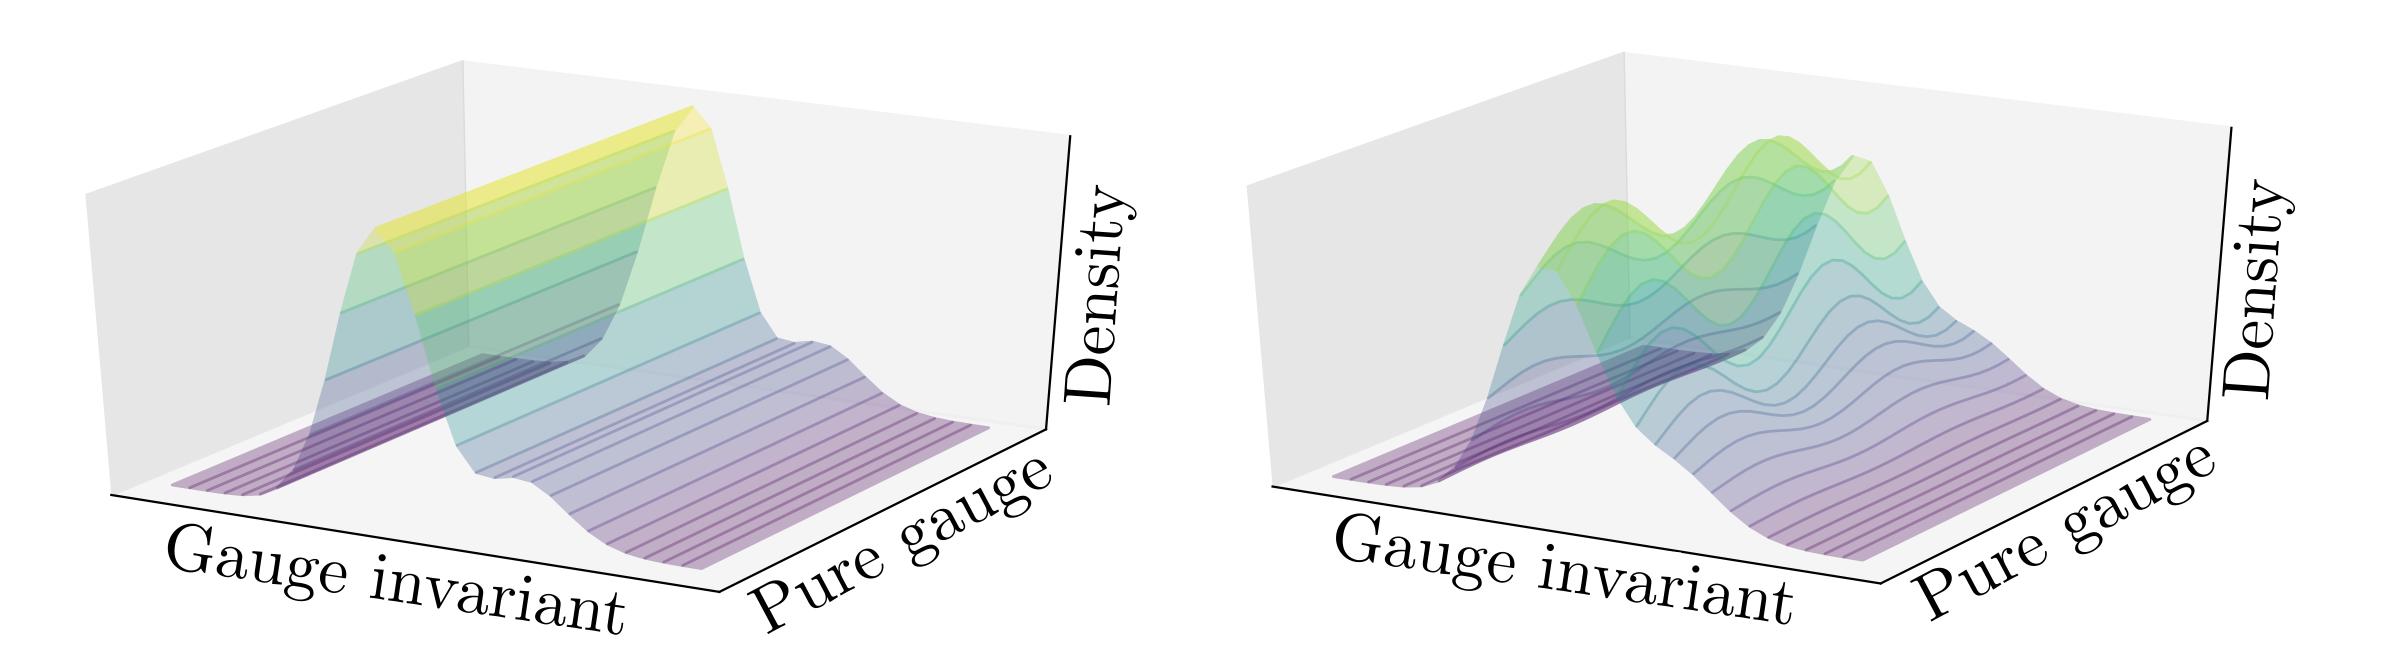


Space-time & Local Gauge Symmetry

The action is invariant to **local gauge transformations**, so the distribution is constant in those directions. **It's a huge product group!** SU(3)(#lattice sites)

Many more pure gauge degrees of freedom than physical ones

We would like to enforce this symmetry in the network, and not have to learn it.



Step 1: Flows on Spheres and Tori

We designed flows on compact manifolds like Spheres and Tori that correspond

to Lie groups:

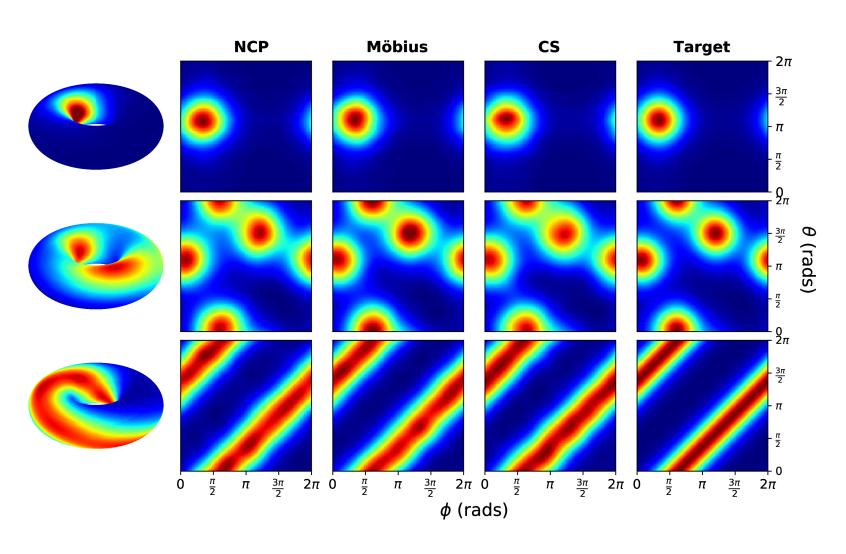


Figure 3. Learned densities on \mathbb{T}^2 using NCP, Möbius and CS flows. Densities shown on the torus are from NCP.

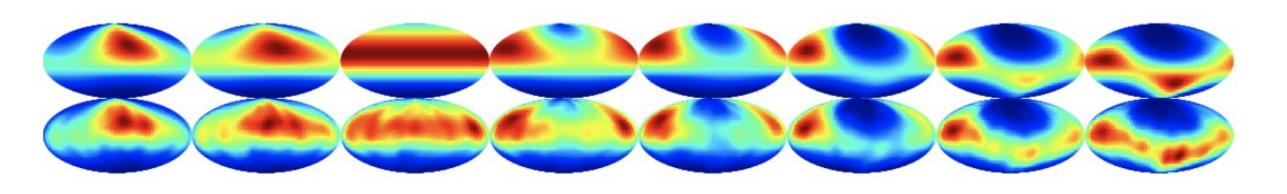
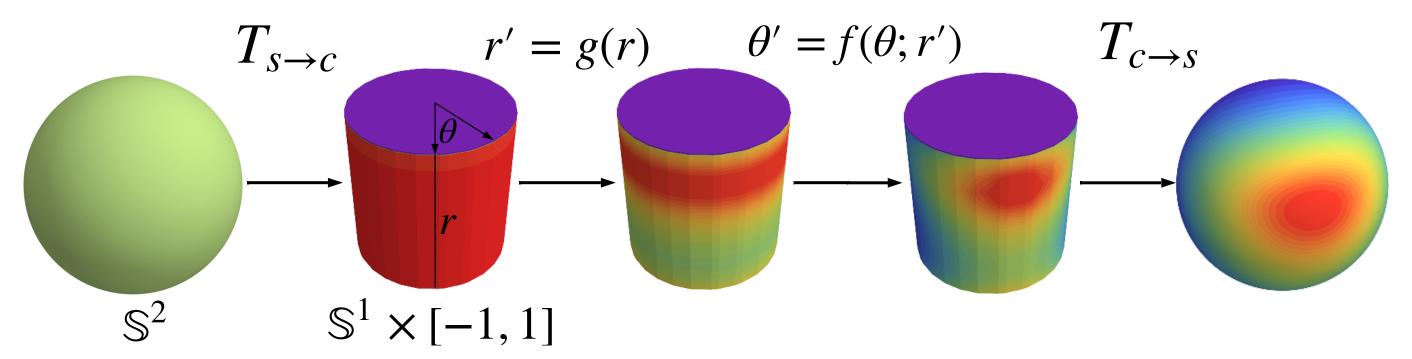
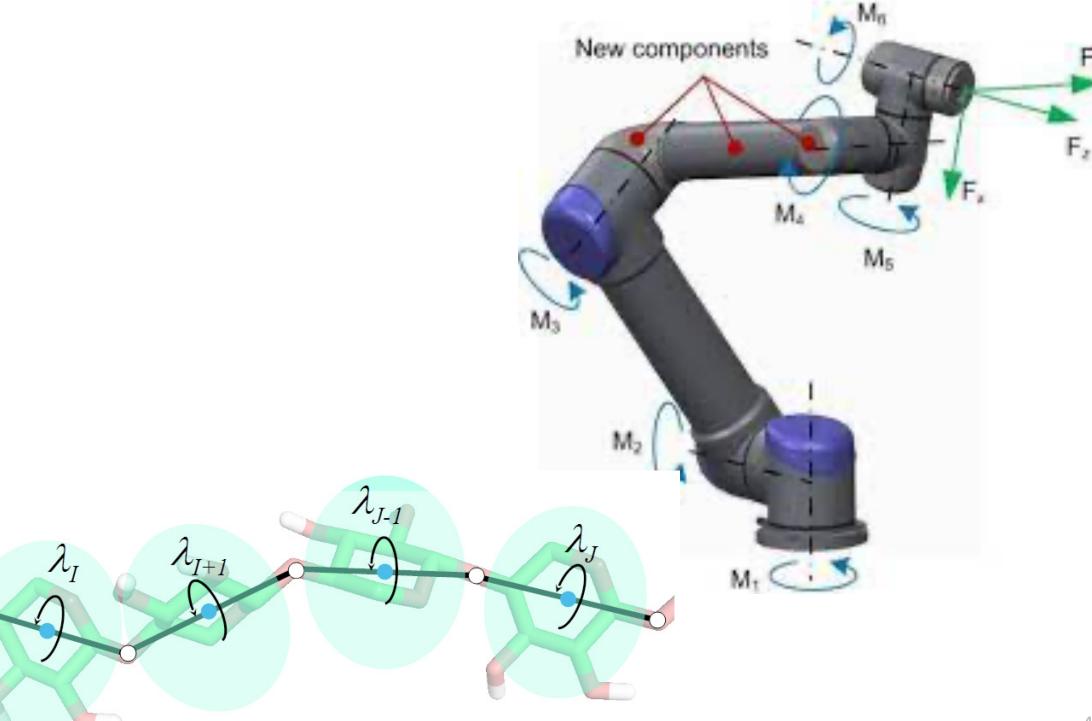


Figure 5. Learned multi-modal density on $SU(2) \equiv \mathbb{S}^3$ using the recursive flow. Each column shows an \mathbb{S}^2 slice of the \mathbb{S}^3 density





Al for Lattice Field Theory

Review ⇒

nature reviews physics

Perspective Check for update

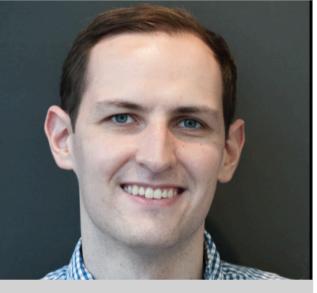
Advances in machine-learning-based sampling motivated by lattice quatum chromodynamics

Kyle Cranmer ® 1, Gurtej Kanwar ® 2, Sébastien Racanière ® 3, Danilo J. Rezende ® 3 & Phiala E. Shanahan ® 4.5 \(\subseteq \)





Abbott



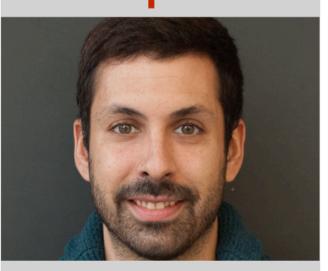
Institute of

Hackett



Massachusetts
Institute of
Technology

Romero-López



Massachusetts
Institute of
Technology

Boyda



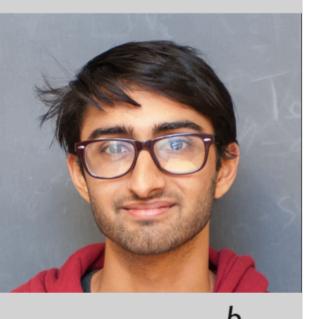


Urban





Kanwar



Botev



DeepMind

Matthews



DeepMind

Rezende





Racanière



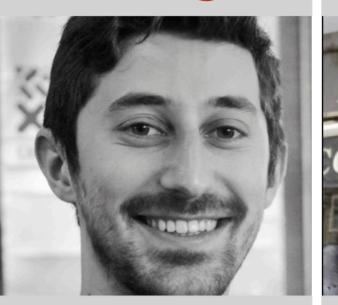


Razavi





Albergo





Cranmer



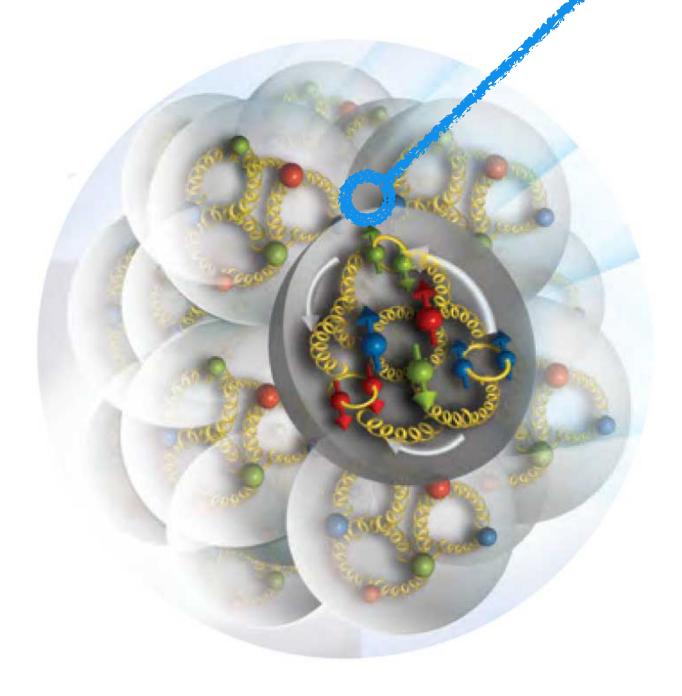


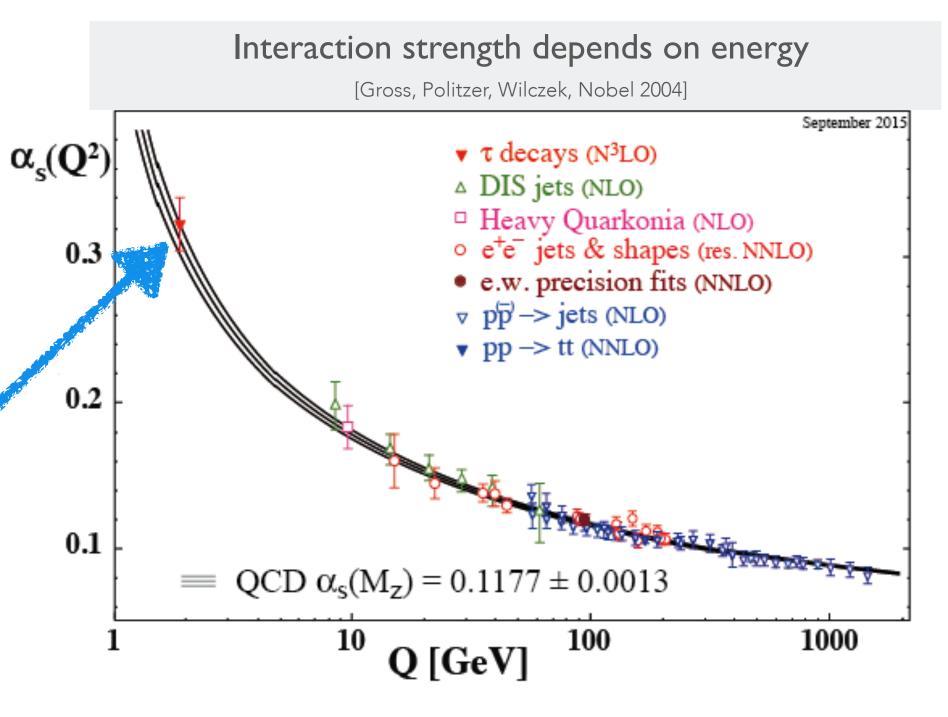
Resource of Inspiration

The strong force: Quantum Chromodynamics (QCD)

QCD is strong at at lowenergies, no small coupling, perturbation theory fails.

Emergent phenomena: protons, pions, etc.

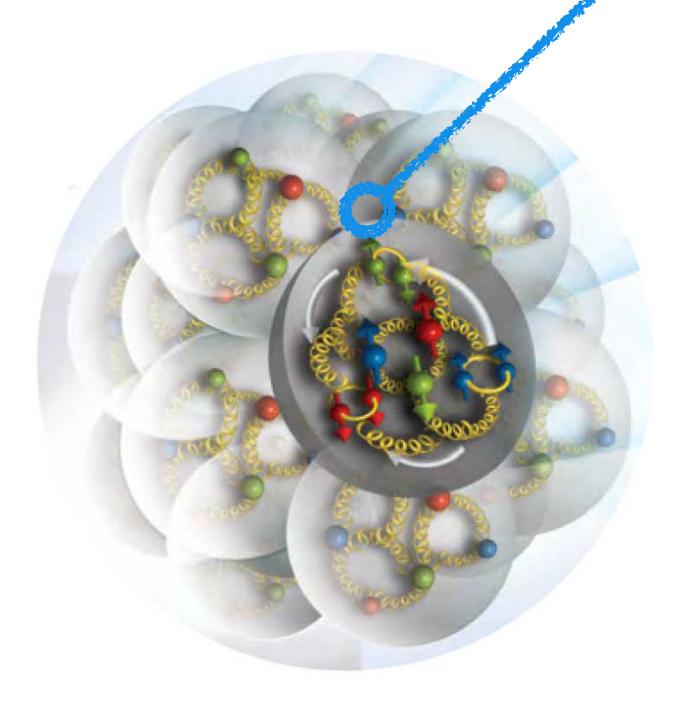


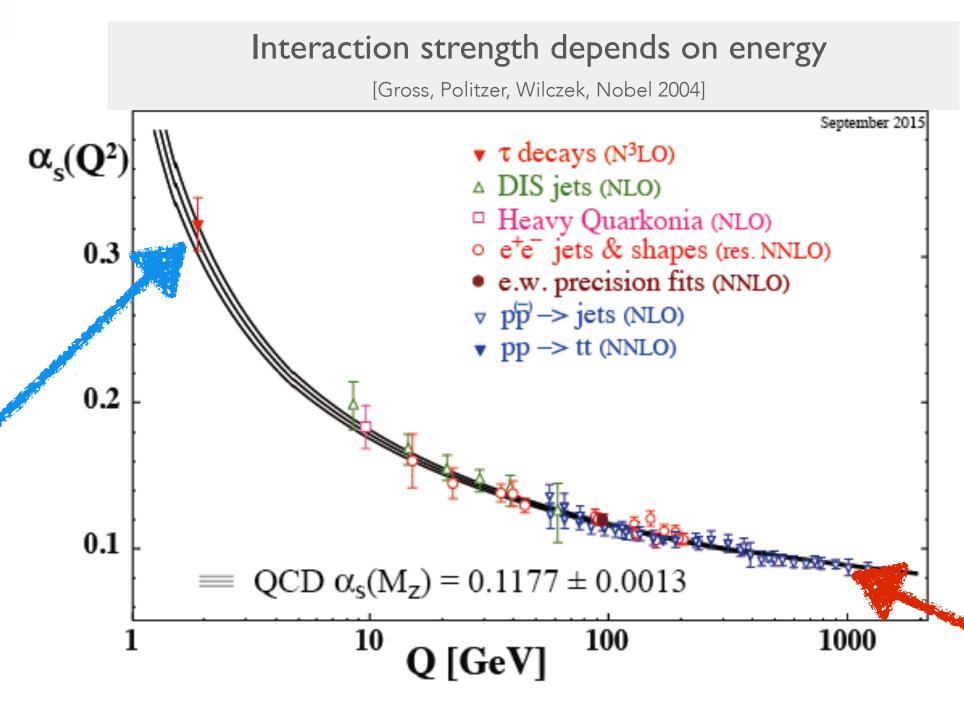


The strong force: Quantum Chromodynamics (QCD)

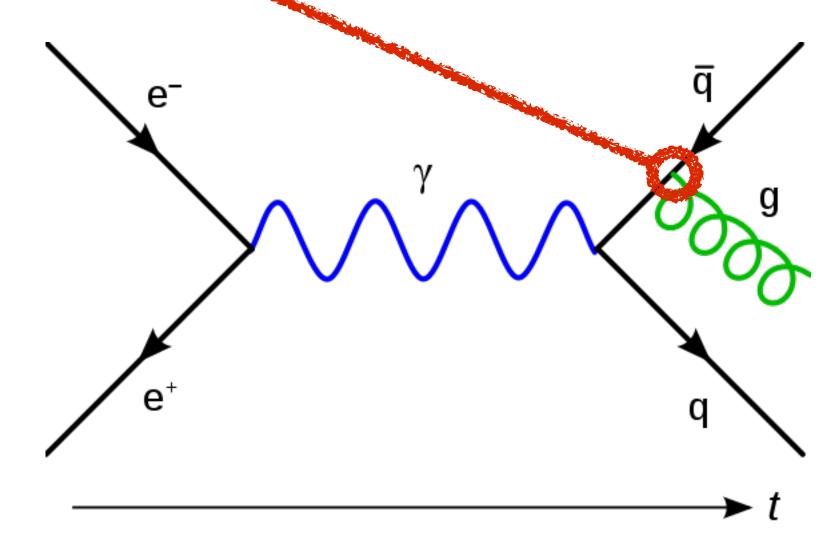
QCD is strong at at lowenergies, no small coupling, perturbation theory fails.

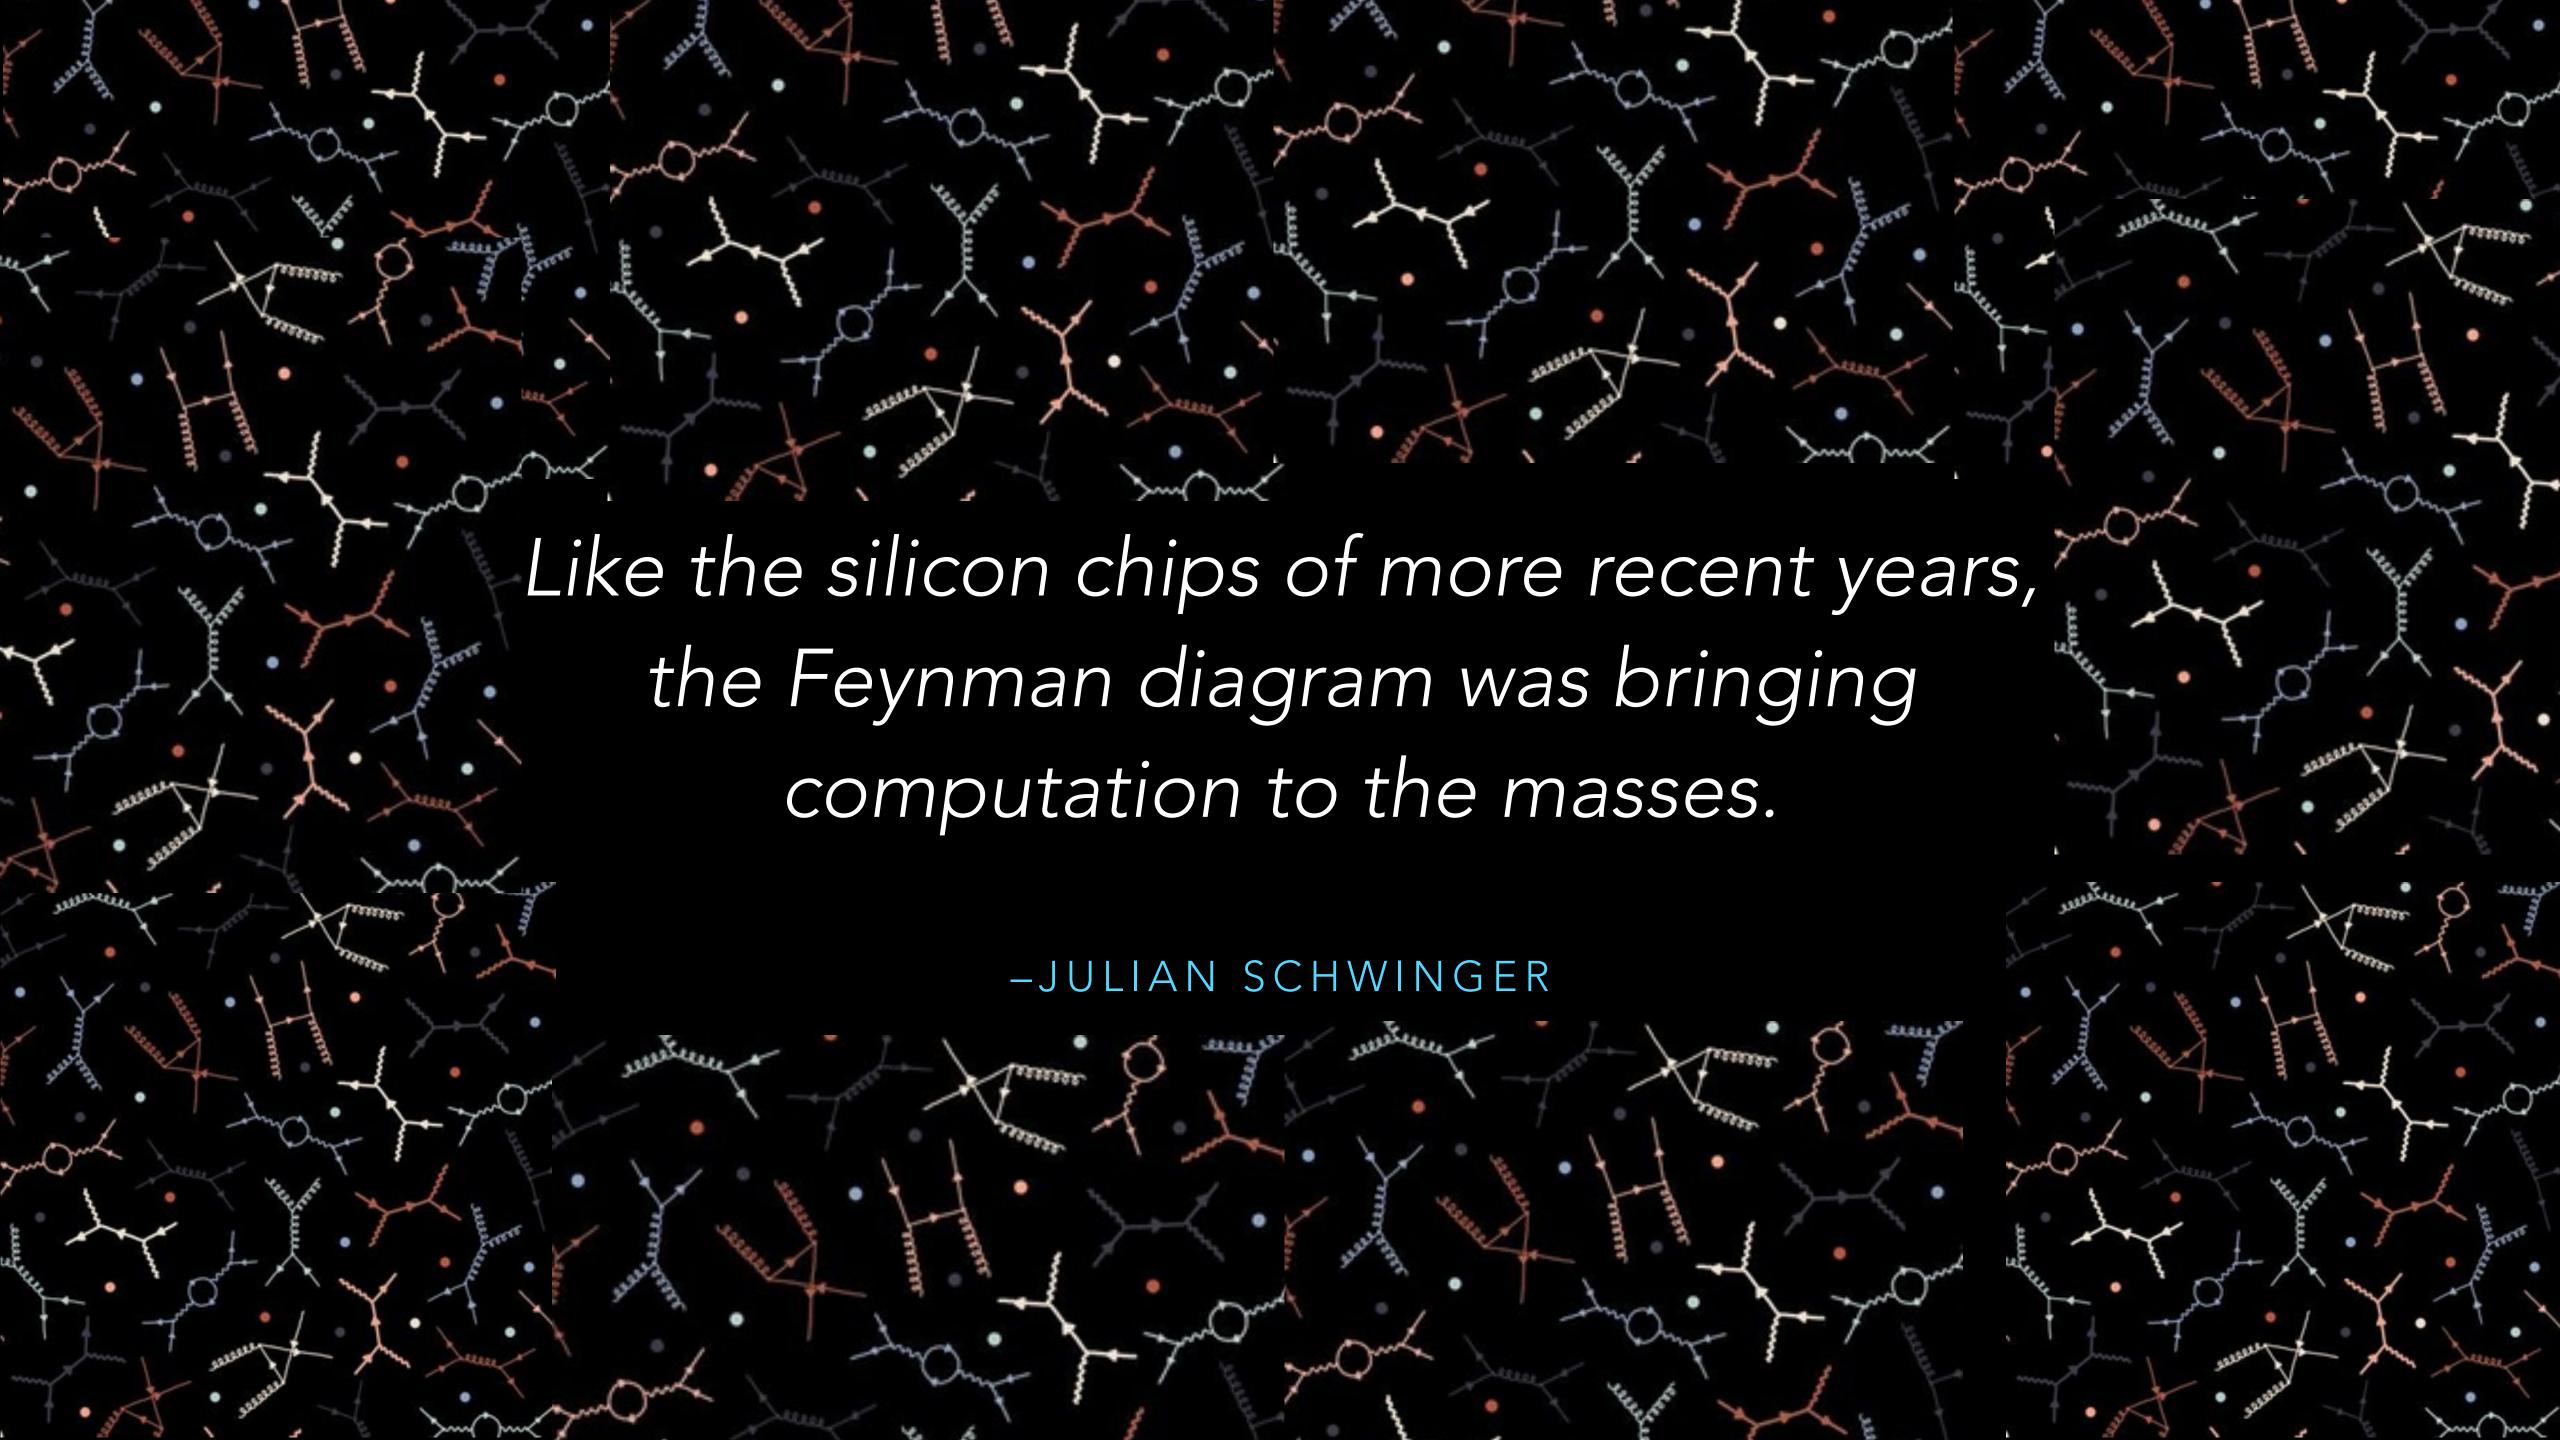
Emergent phenomena: protons, pions, etc.





QCD is weak at at highenergies, small coupling, perturbation theory works





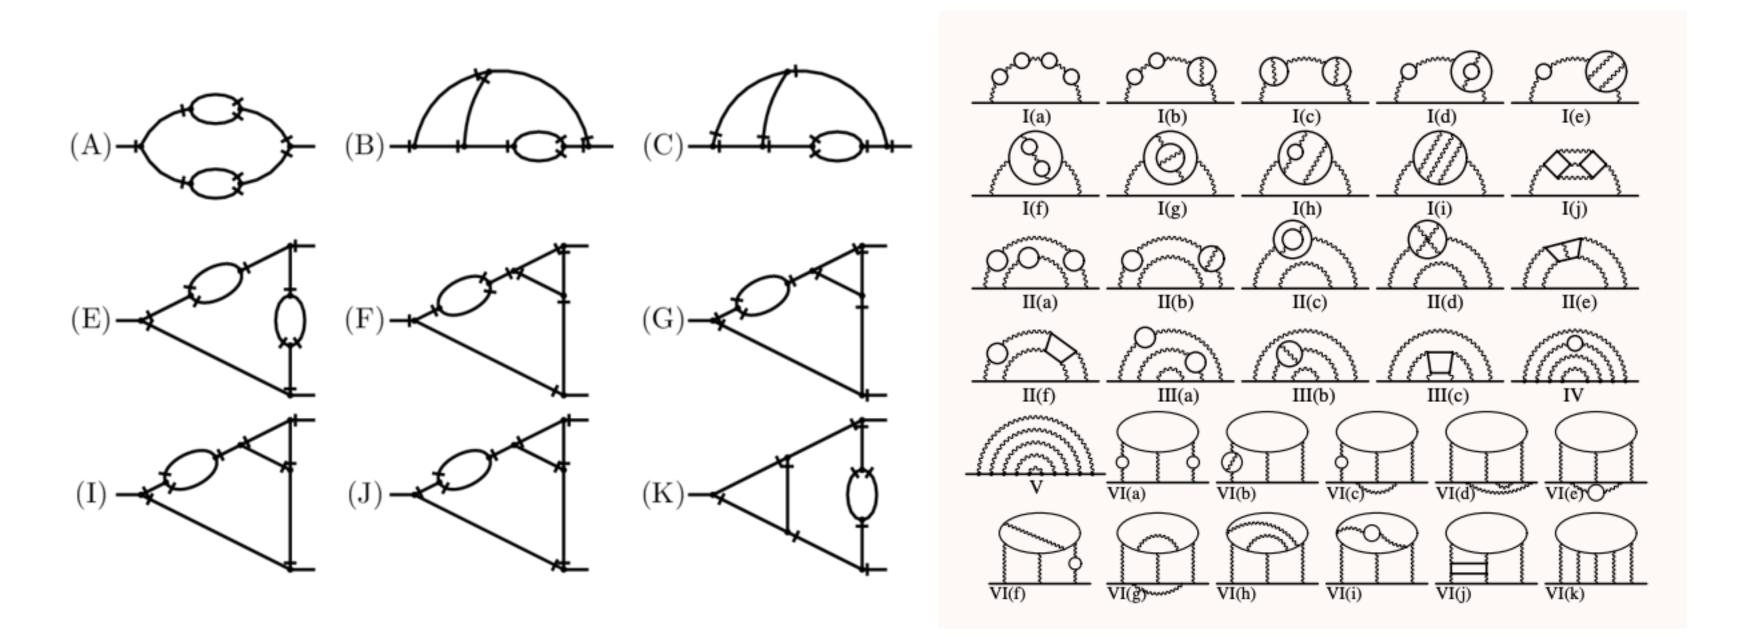
Feynman diagrams with loops

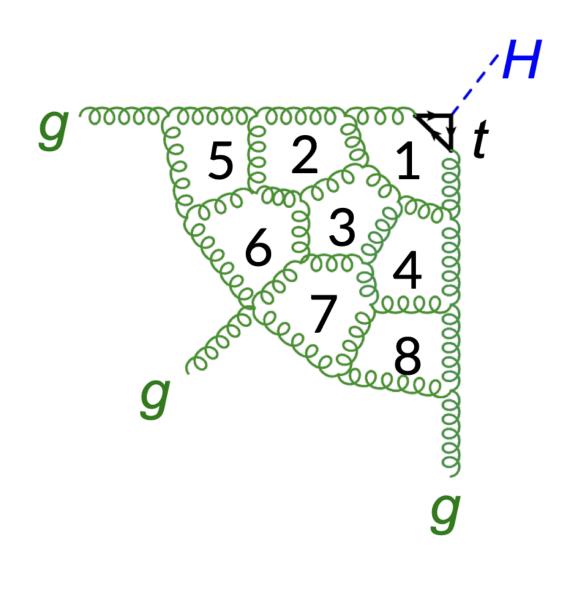
Galileo Galilei Medal 2023 Zvi Bern, Lance Dixon e David Kosower

More precise calculations have more loops

- But the number of diagrams grows combinatorially with the number of loops
- Feynman diagrams become a poor way to organize the calculation

New **bootstrap** approach emerged that leverages analytical properties of amplitudes. Properties are so constraining, they define a unique solution





Bootstrap

In this bootstrap approach the L-loop amplitude can be expressed as a sum of **terms** with an **integer coefficient** and a **word** composed of 2L **letters**

• The 6 letters $\{a, b, c, d, e, f\}$ encode the kinematics of the collision

Bootstrap

In this bootstrap approach the L-loop amplitude can be expressed as a sum of **terms** with an **integer coefficient** and a **word** composed of 2L **letters**

• The 6 letters $\{a, b, c, d, e, f\}$ encode the kinematics of the collision

For example, in a particular theory called $\mathcal{N}=4$ Super Yang-Mills theory, the answer at 2-loops for a particular interaction is:

$$S[F_3^{(2)}] = +8bddd + 8ceee + 8afff + 8bfff + 8cddd + 8aeee$$
$$+16bbbd + 16ccce + 16bbbf + 16aaaf + 16cccd + 16aaae$$

• Of the $6^4 = 1296$ possible terms, most are 0. Sparse, lots of structure!

Bootstrap

In this bootstrap approach the L-loop amplitude can be expressed as a sum of **terms** with an **integer coefficient** and a **word** composed of 2L **letters**

• The 6 letters $\{a, b, c, d, e, f\}$ encode the kinematics of the collision

For example, in a particular theory called $\mathcal{N}=4$ Super Yang-Mills theory, the answer at 2-loops for a particular interaction is:

• P1

$$S[F_3^{(2)}] = +8bddd + 8ceee + 8afff + 8bfff + 8cddd + 8aeee$$
$$+16bbbd + 16ccce + 16bbbf + 16aaaf + 16cccd + 16aaae$$

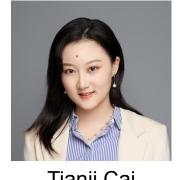
• Of the $6^4 = 1296$ possible terms, most are 0. Sparse, lots of structure!

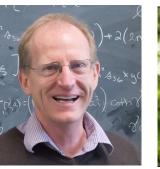
The solution space is growing exponentially! Hard to find the answer

=	${\rm loop\ order}\ L$	1	2	3	4	5	6	7	8
_	terms in $S[F_3^{(L)}]$	6	12	636	11,208	263,880	4.9×10^{6}	9.3×10^{7}	1.67×10^{9}

Can Al help?









Matthias Wilhelm Niklas Nolte





Similar to an NP-complete problem:

• the answer is hard to find, but easy to check.

We don't need the model to be provably correct, we just need it to be good at guessing because we can get a certificate of correctness

• The problem is inherently discrete, so transformers are a natural choice



TRANSFORMERS FOR SCATTERING AMPLITUDES

Garrett W. Merz †, Tianji Cai *, François Charton § Niklas Nolte §, Matthias Wilhelm ‡, Kyle Cranmer †, Lance Dixon *





Can Transformers learn N=4 Super Yang Mills Theory?

• 1, 1, 2, 3, 5, 8, 13, 21, ...

No i.i.d. data, only one example.

What is the recursion rule?

•
$$x_n = x_{n-1} + x_{n-2}$$

Our problem is much harder:

 each element of the sequence is a "tensor" growing in dimension

$$\mathbb{Z}^{6^2} \to \mathbb{Z}^{6^4} \to \mathbb{Z}^{6^6} \to \mathbb{Z}^{6^8} \to \mathbb{Z}^{6^{2L}}$$



Transformers can discover recurrence relations from sequences (aka IQ tests). New paper on symbolic regression, with @stephanedascoli @pa_kamienny and @GuillaumeLample

Guillaume Lample @ #NeurIPS2023 @GuillaumeLample · Jan 14, 2022 Deep Symbolic Regression for Recurrent Sequences -arxiv.org/abs/2201.04600 We show that transformers are great at predicting symbolic functions from values, and can predict the recurrence relation of sequences better than Mathematica. You can try it here: bit.ly/3niE5FS 200 -0.0

100

Building Community

The ML4Jets Workshops







2017 2020

See agendas here: https://indico.cern.ch/event/1253794/

Hammers & Nails





NeurlPS workshops



Machine Learning and the Physical Sciences

Workshop at the 37th conference on Neural Information Processing Systems (NeurIPS)

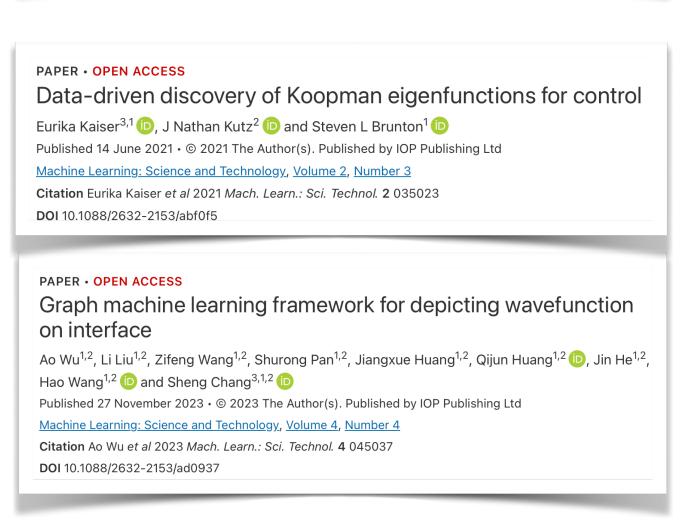
December 15, 2023

A journal to consider

Machine Learning: Science and Technology is a multidisciplinary open access journal that bridges the application of machine learning across the sciences with advances in machine learning methods and theory as motivated by physical insights. Disclosure: I'm the current Editor-in-Chief

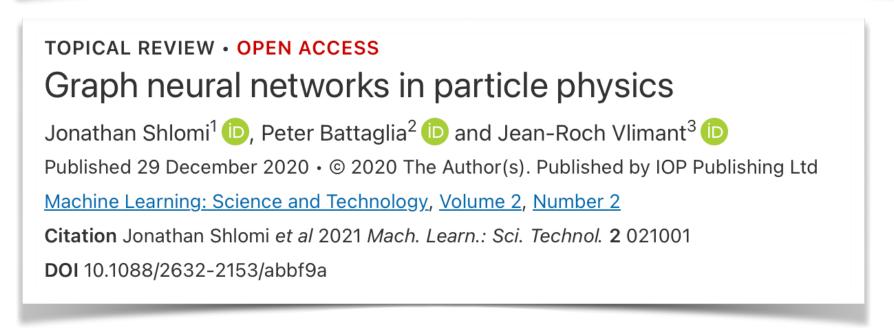


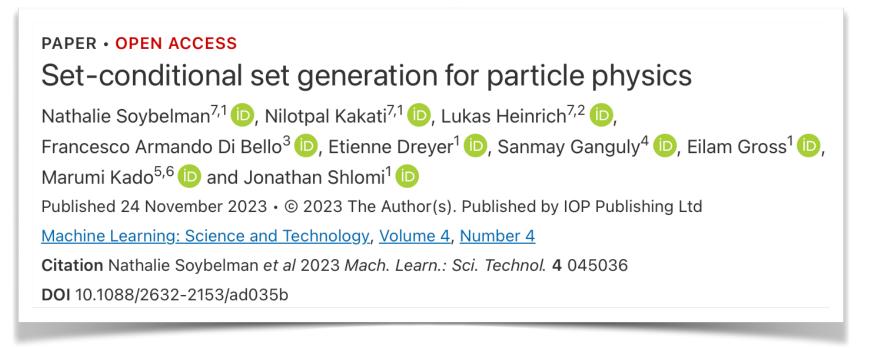






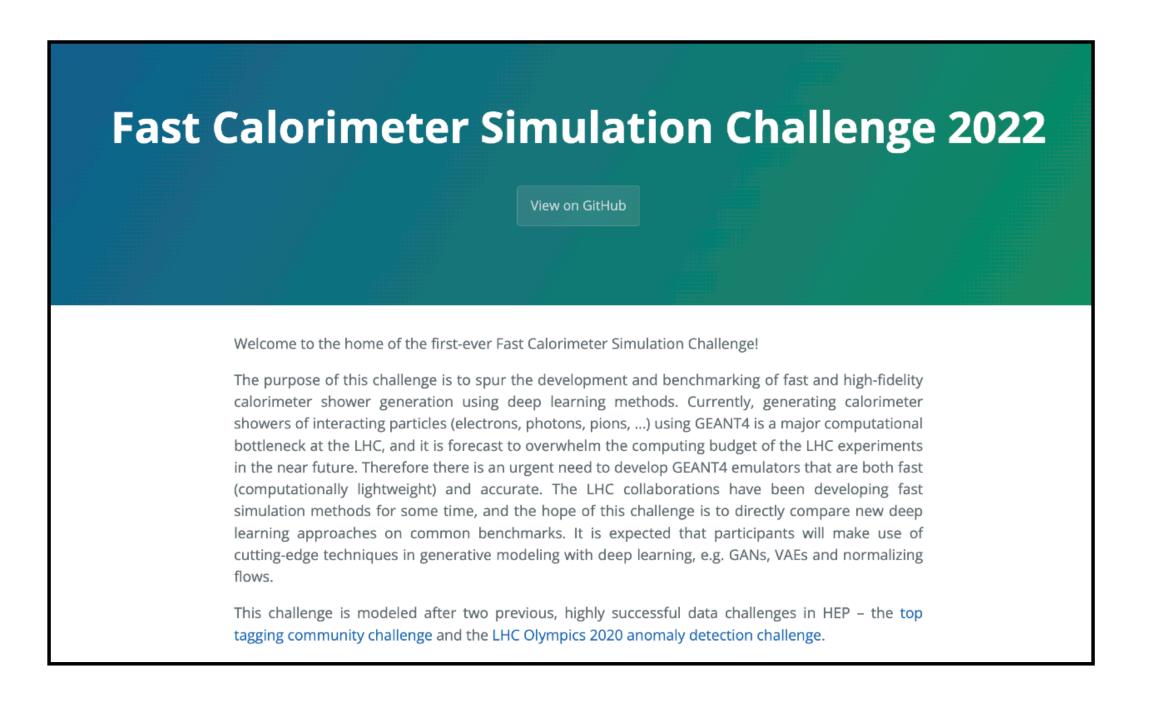
PAPER • OPEN ACCESS Rediscovering orbital mechanics with machine learning Pablo Lemos 9,1,2 , Niall Jeffrey 2,3,4, Miles Cranmer 5, Shirley Ho 5,6,7,8 and Peter Battaglia Published 9 October 2023 • © 2023 The Author(s). Published by IOP Publishing Ltd Machine Learning: Science and Technology, Volume 4, Number 4 Citation Pablo Lemos et al 2023 Mach. Learn.: Sci. Technol. 4 045002 DOI 10.1088/2632-2153/acfa63

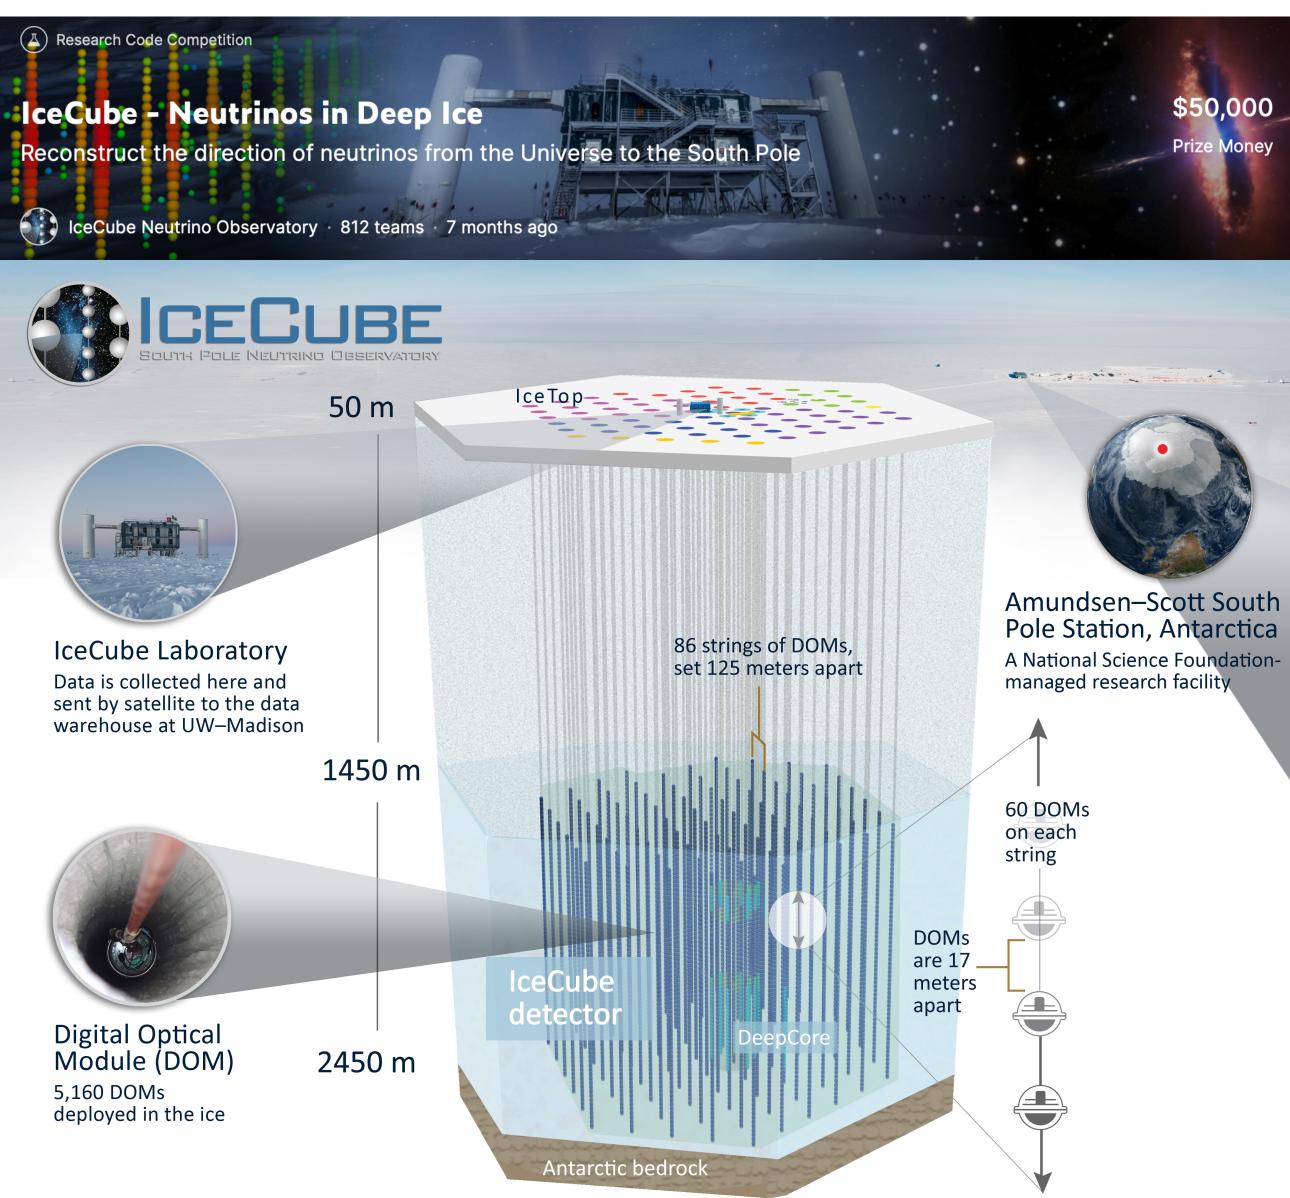




Challenges, Benchmarks, and Open Data







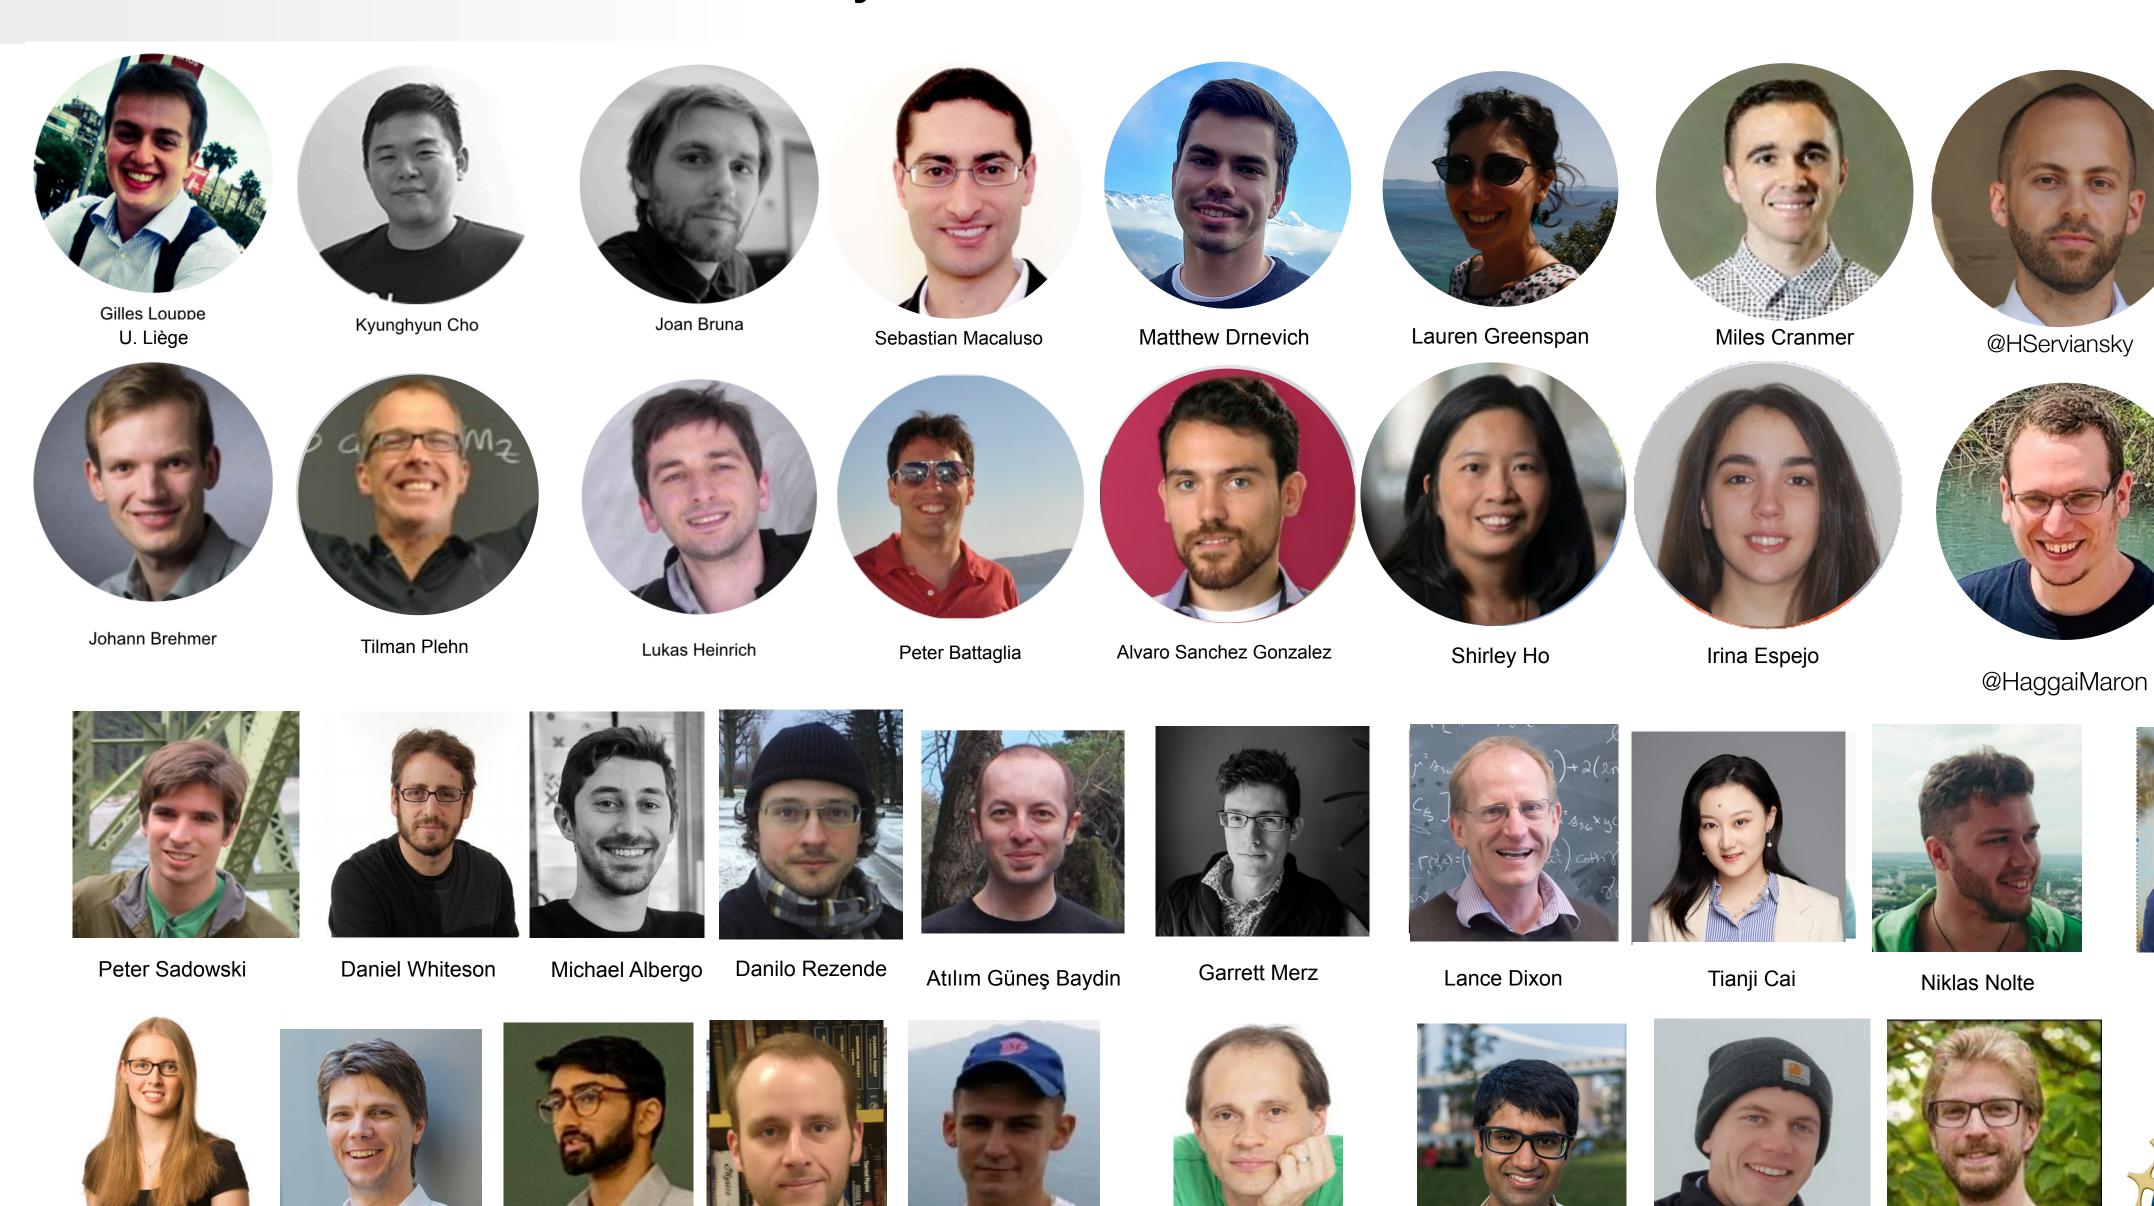
Collaborators + Many More

William Detmold

Dan Hackett

Gurtej Kanwar

Phiala Shanahan



Denis Boyda

Sébastien Racanière

Patrick Rieck

Matthias Wilhelm

Siddharth Mishra-Sharma

65

Nimrod Segol

Yaron Lipman

Francois Charton

Fundamental Physics has a compelling story to tell and an important role in Al4Science.

... and it is critical that we tell it!

EuCAIF can help

Thank you!

Questions?