Particle Physics and the AI revolution

Juan Rojo, VU Amsterdam & Nikhef

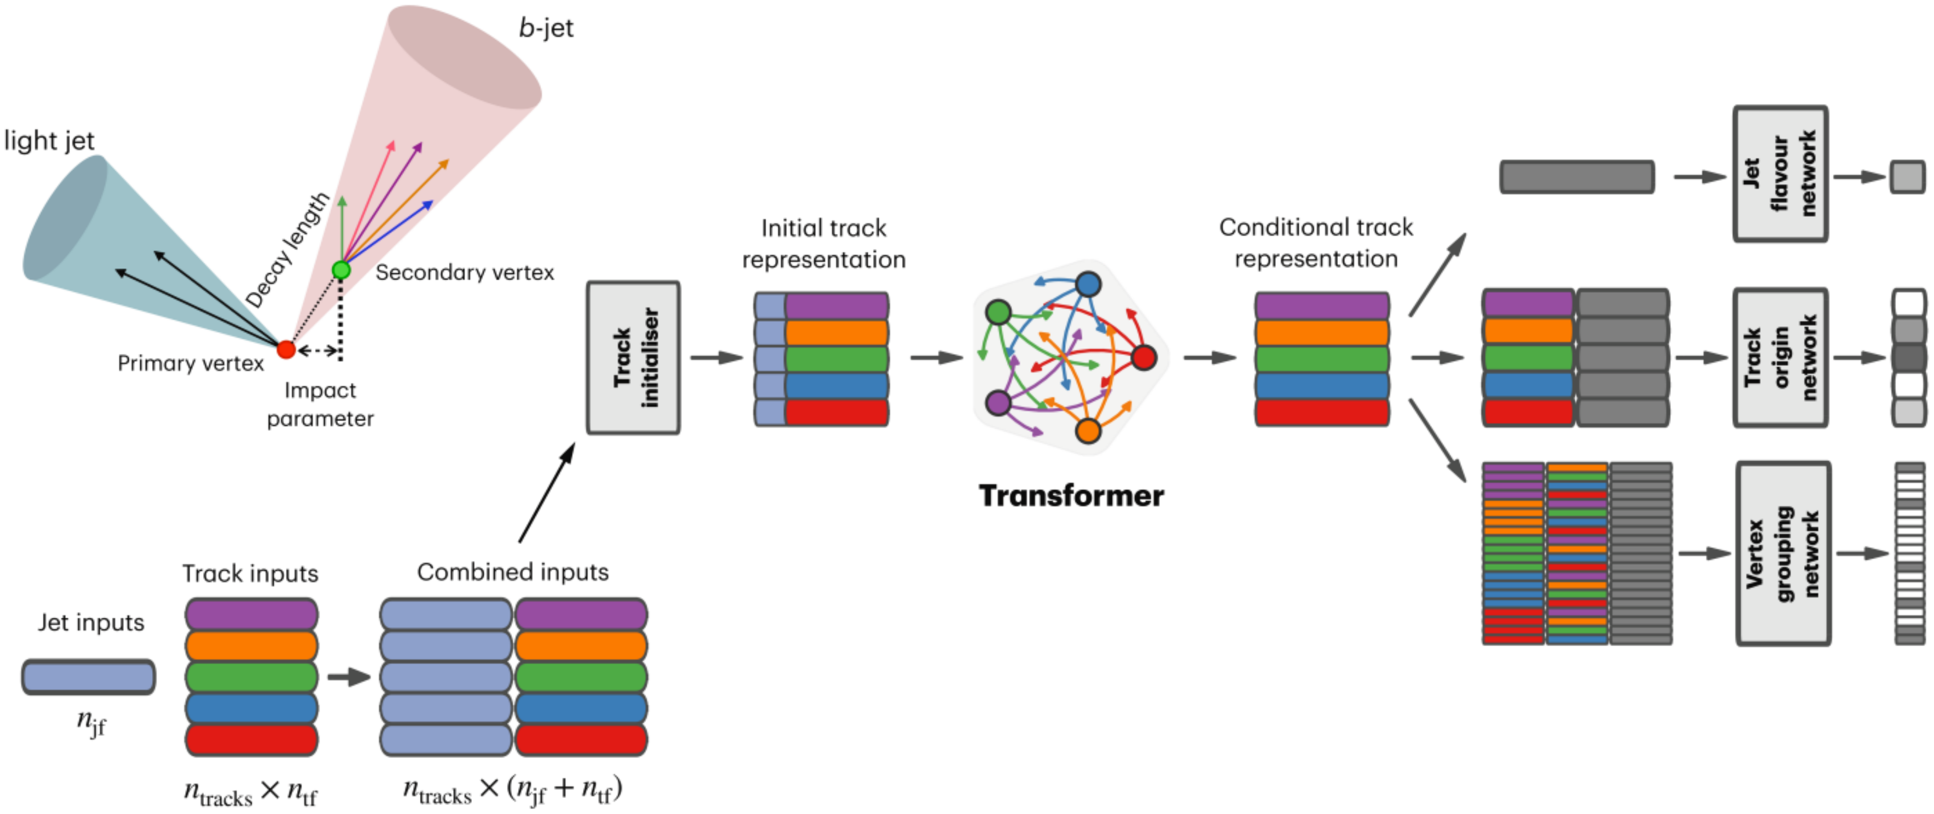

Institut de Fisica d'Altes Energies (IFAE) Colloquium,

March 23rd 2026

Journeys into the Zeptospace

Journeys into the Zeptospace

lenght units $\rightarrow L = 10^{-10}$ m (Bohr's radius, size of H atom)

energy/mass units $\rightarrow E = 10^{9}$ eV $= 1$ GeV (mass of H atom, $E=mc^2$)

long distances
& low energies

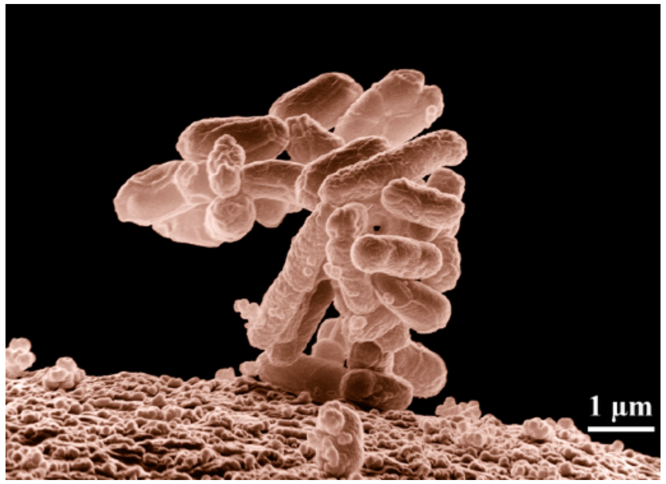

bacteria (10^4 *L*)

short distances
& large energies

Journeys into the Zeptospace

lenght units → $L = 10^{-10}$ m (Bohr's radius, size of H atom)

energy/mass units → $E = 10^{9}$ eV $= 1$ GeV (mass of H atom, $E=mc^2$)

long distances
& low energies

virus (10^3 *L*)

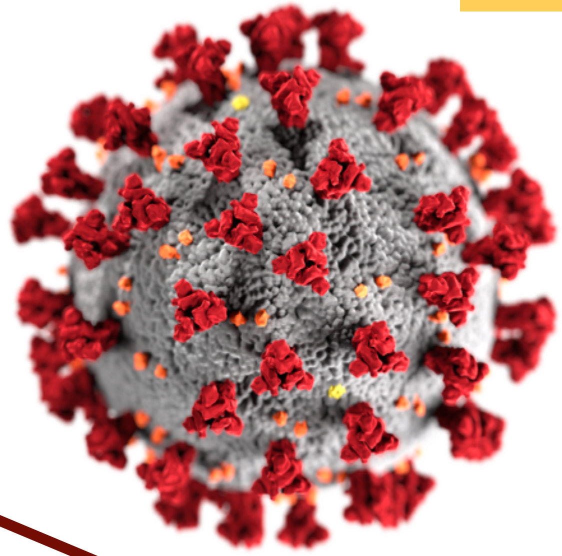

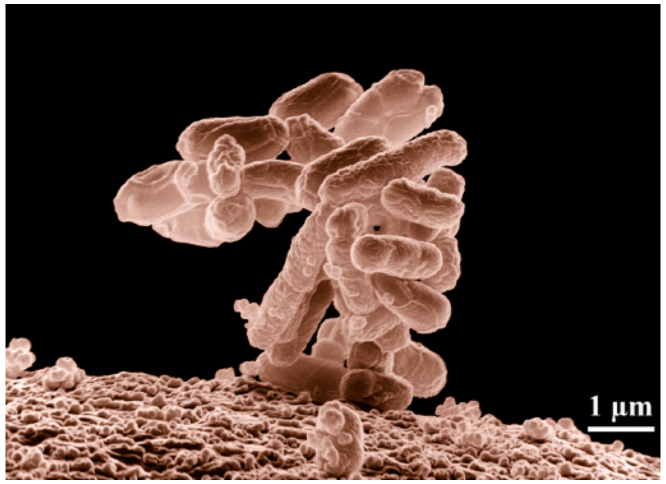

bacteria (10^4 *L*)

short distances
& large energies

Journeys into the Zeptospace

$$\text{lenght units} \rightarrow L = 10^{-10}\ \text{m}$$ (Bohr's radius, size of H atom)

$$\text{energy/mass units} \rightarrow E = 10^{9}\ \text{eV} = 1\ \text{GeV}$$ (mass of H atom, $E=mc^2$)

long distances
& low energies

virus (10^3 *L*)

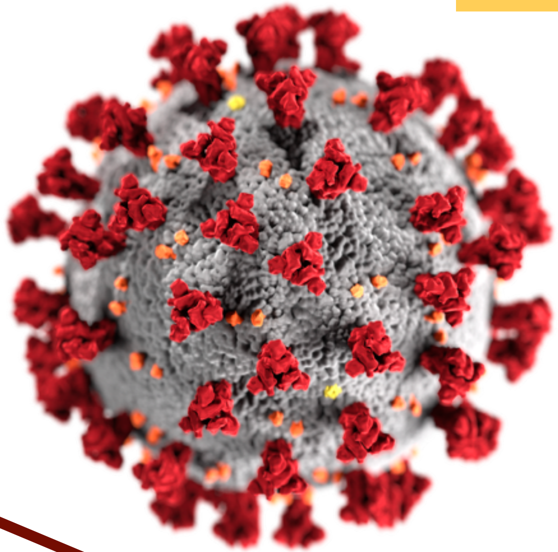

atoms (10^0 *L*, 10^0 *E*)

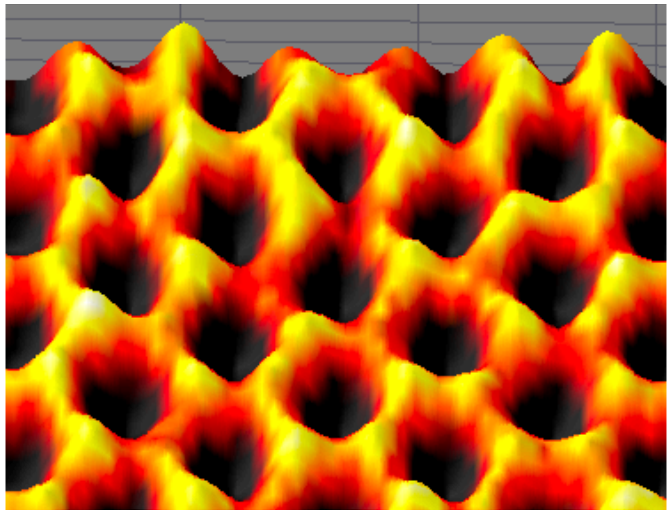

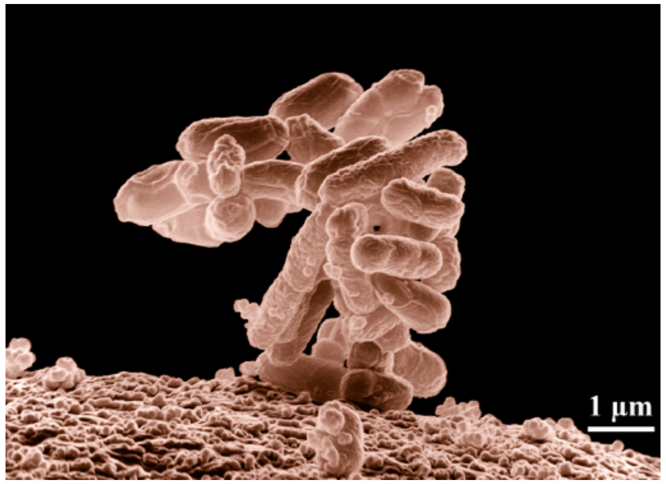

bacteria (10^4 *L*)

short distances
& large energies

Journeys into the Zeptospace

$\text{lenght units} \rightarrow L = 10^{-10}\ \text{m}$ (Bohr’s radius, size of H atom)

$\text{energy/mass units} \rightarrow E = 10^{9}\ \text{eV} = 1\ \text{GeV}$ (mass of H atom, $E=mc^2$)

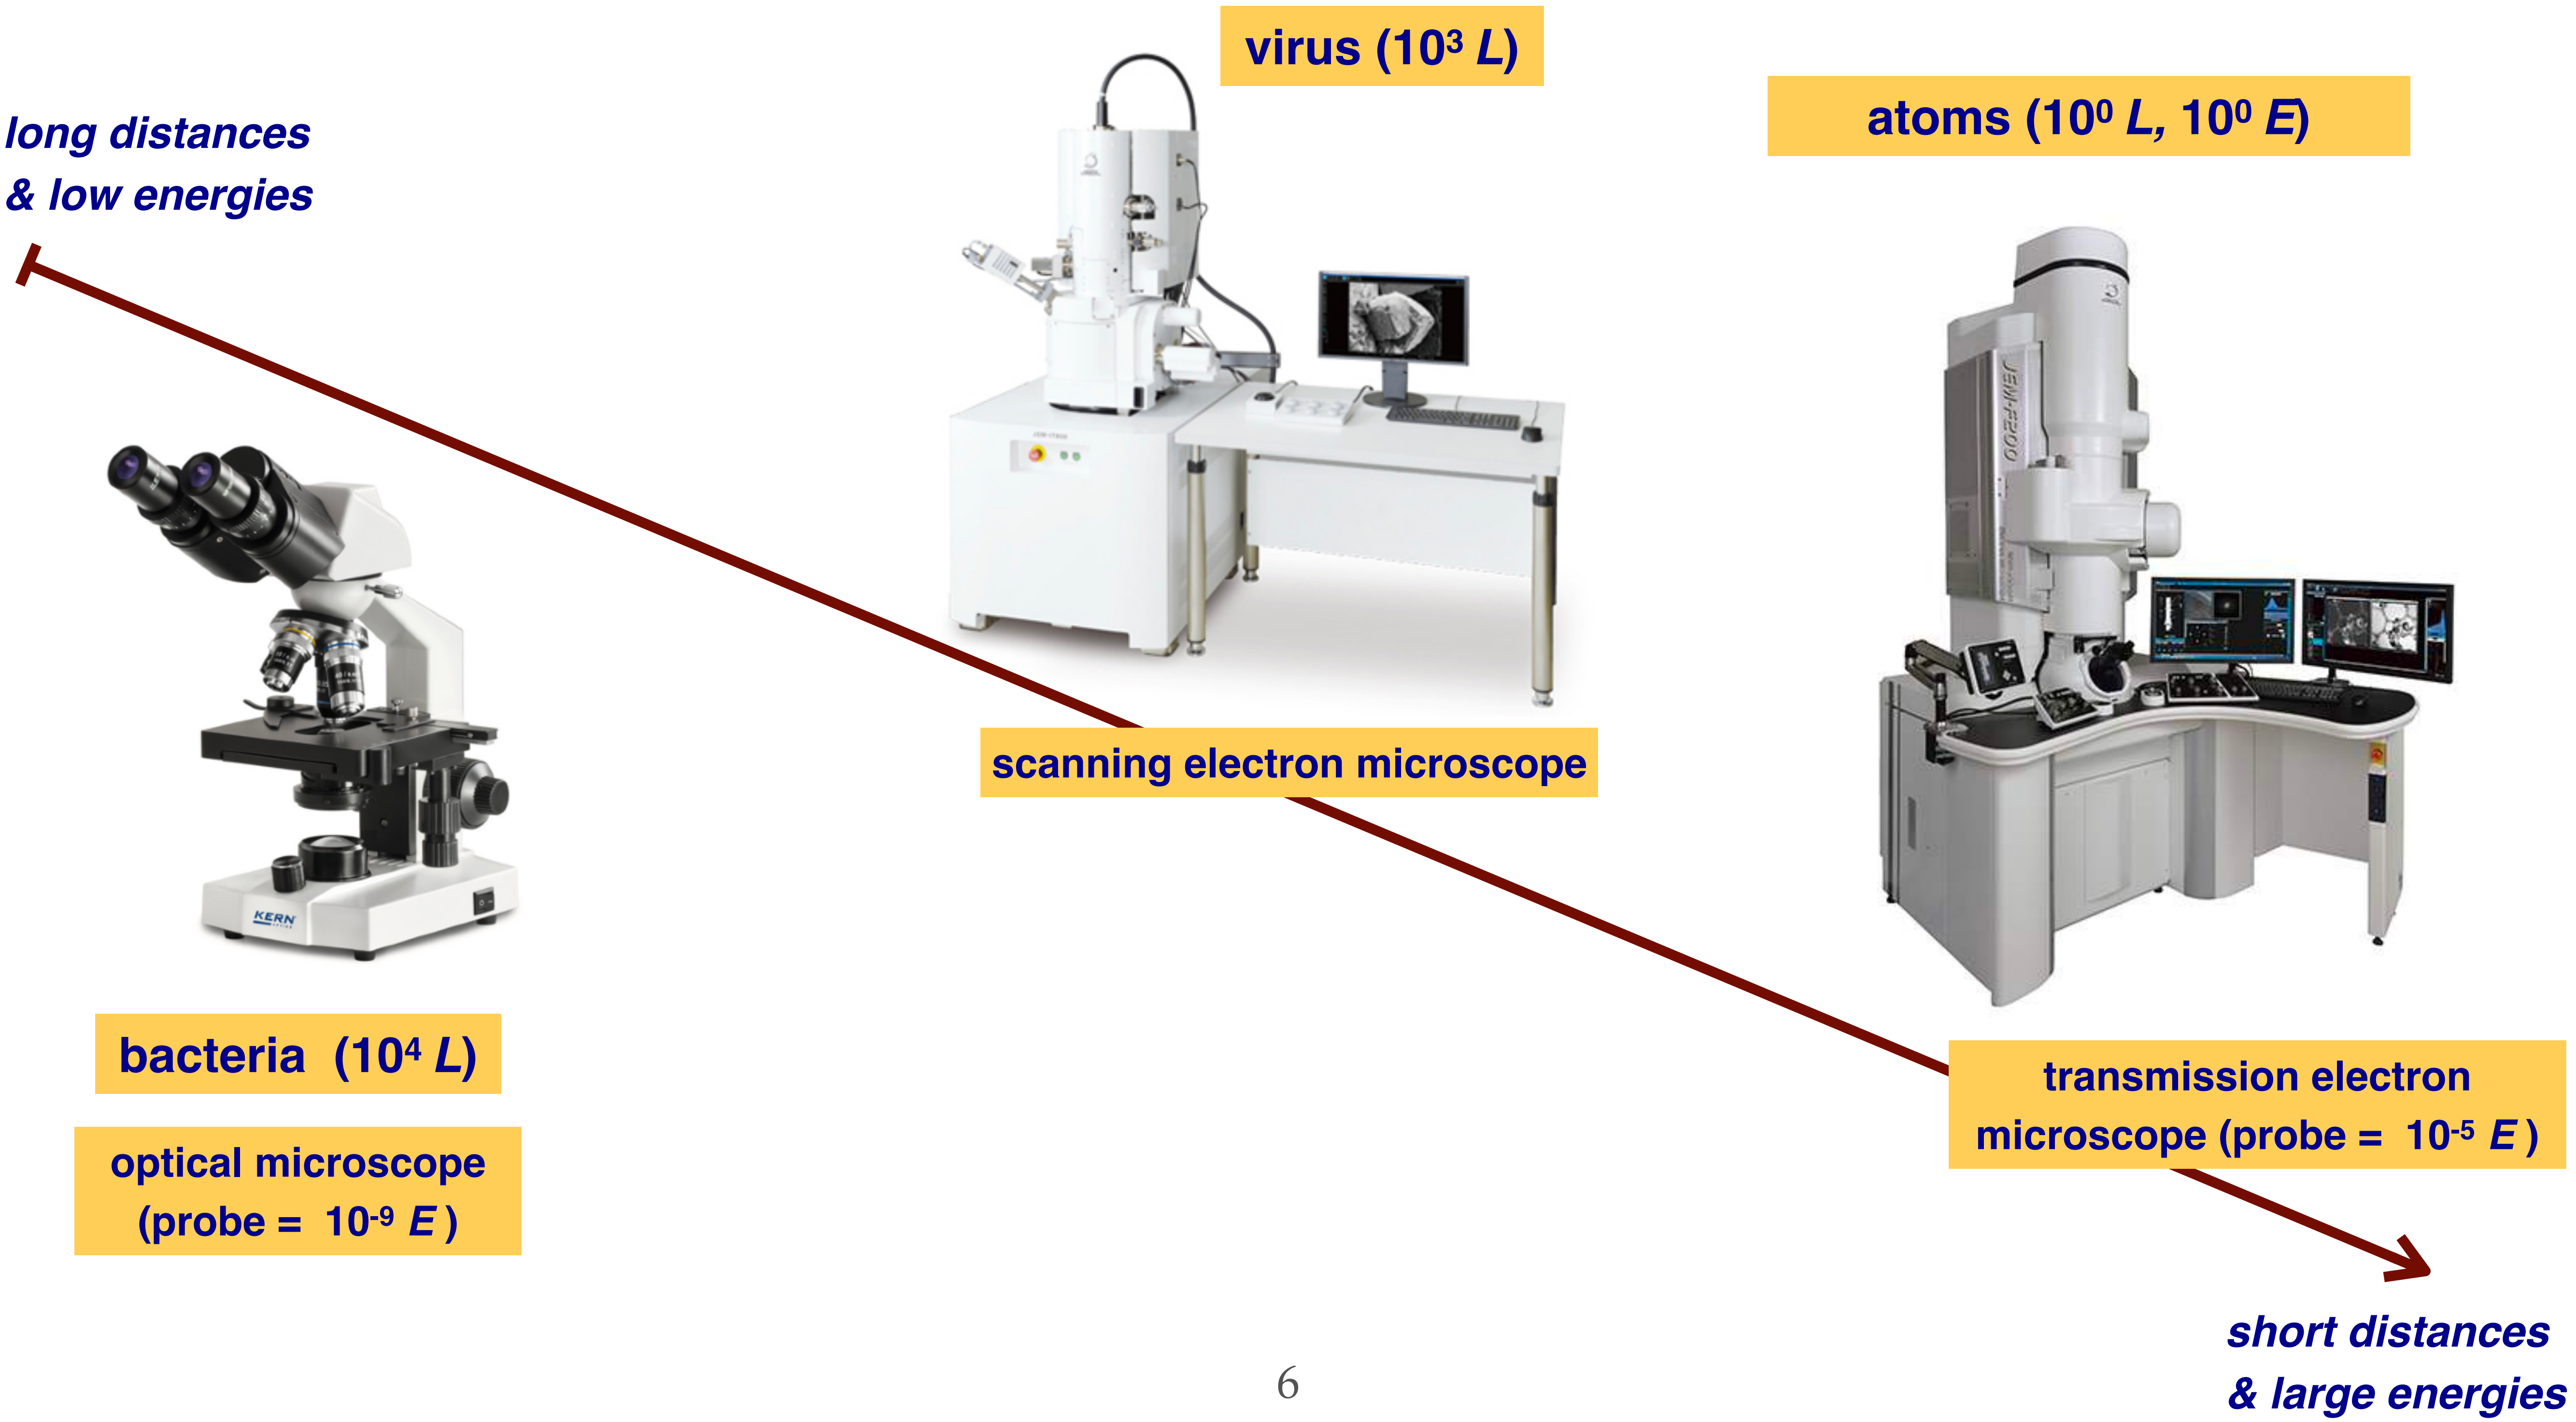

long distances & low energies

bacteria (10^4 *L*)

optical microscope (probe = 10^{-9} *E*)

virus (10^3 *L*)

scanning electron microscope

atoms (10^0 *L*, 10^0 *E*)

transmission electron microscope (probe = 10^{-5} *E*)

short distances & large energies

Journeys into the Zeptospace

lenght units → $L = 10^{-10}$ m (Bohr's radius, size of H atom)

energy/mass units → $E = 10^{9}$ eV = 1 GeV (mass of H atom, $E=mc^2$)

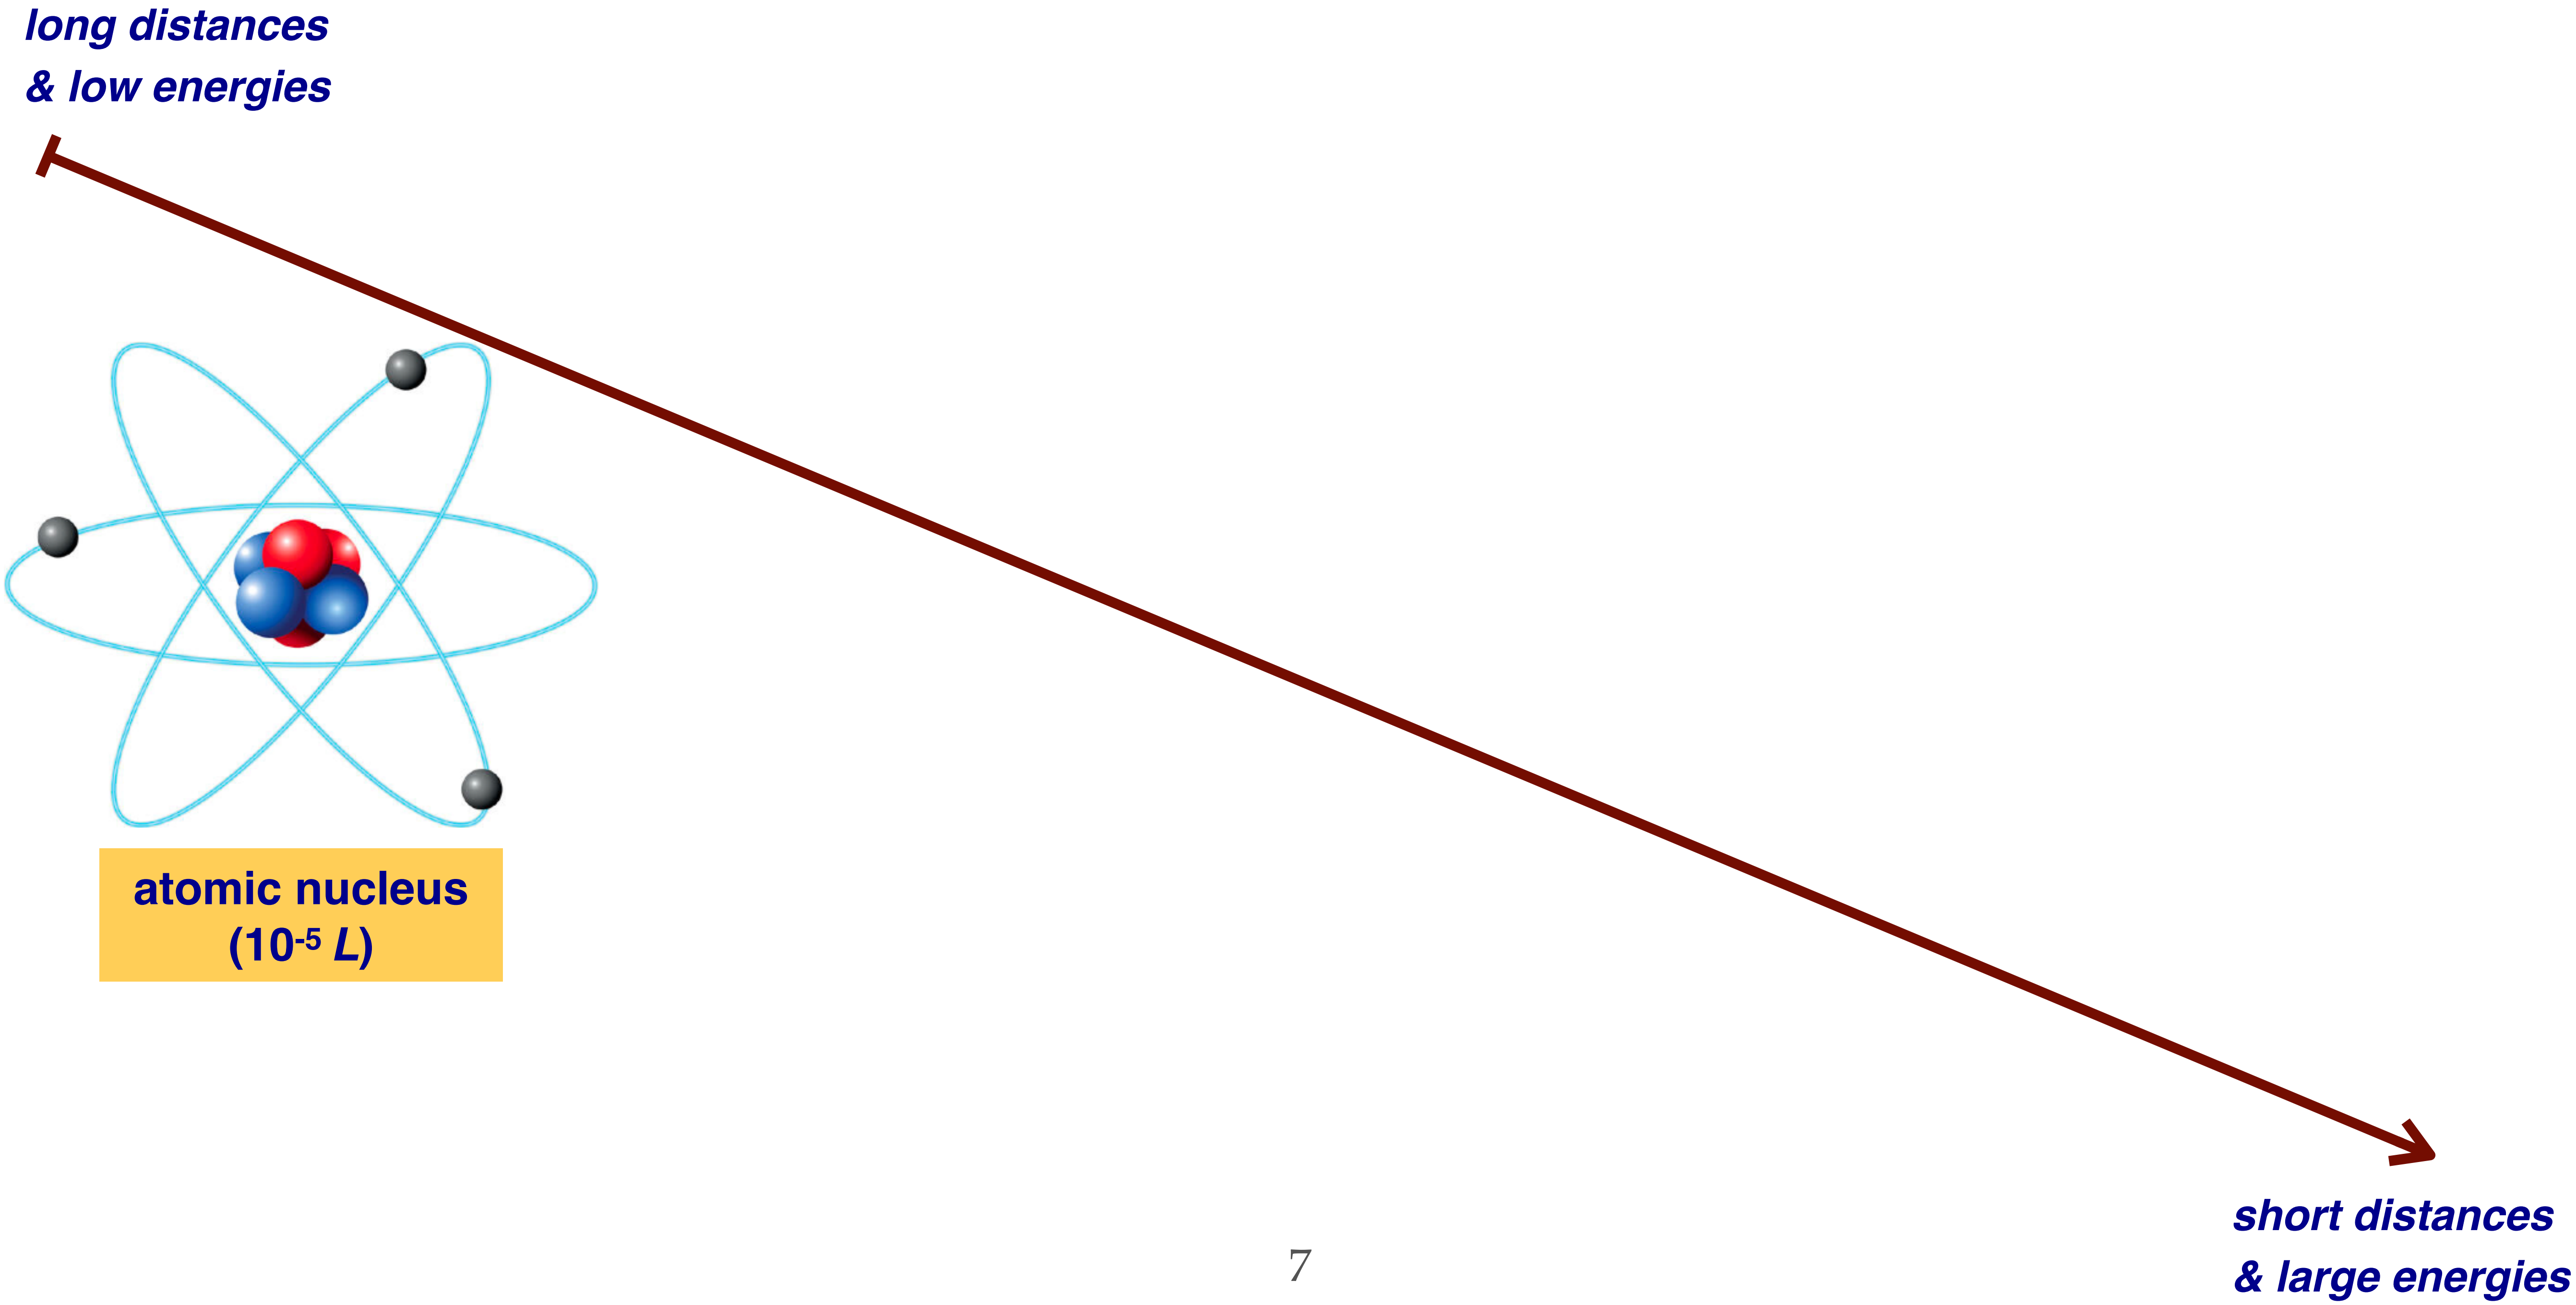

Journeys into the Zeptospace

lenght units $\rightarrow L = 10^{-10}$ m (Bohr's radius, size of H atom)

energy/mass units $\rightarrow E = 10^{9}$ eV $= 1$ GeV (mass of H atom, $E=mc^2$)

quark and gluon structure of protons (10^{-6} *L*)

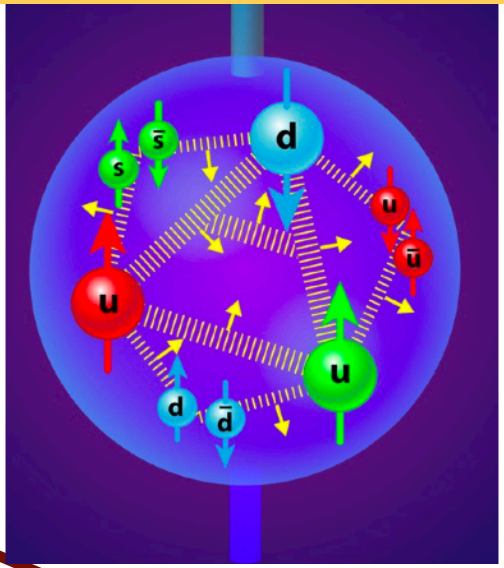

long distances & low energies

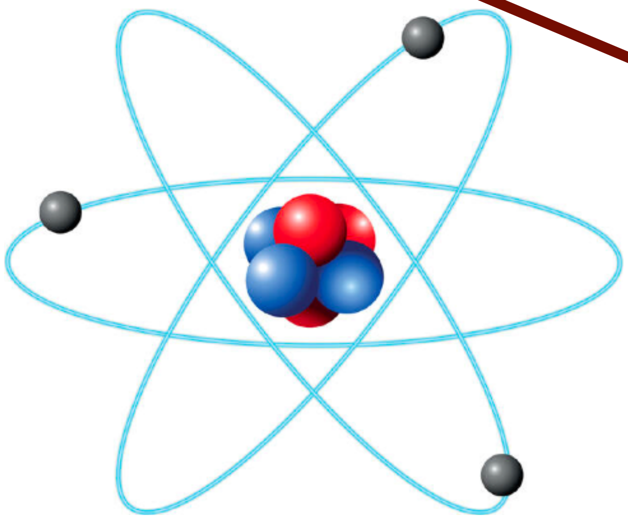

atomic nucleus (10^{-5} *L*)

short distances & large energies

Journeys into the Zeptospace

$\text{lenght units} \rightarrow L = 10^{-10}\ \text{m}$ (Bohr's radius, size of H atom)

$\text{energy/mass units} \rightarrow E = 10^{9}\ \text{eV} = 1\ \text{GeV}$ (mass of H atom, $E=mc^2$)

long distances
& low energies

quark and gluon structure of protons (10^{-6} *L*)

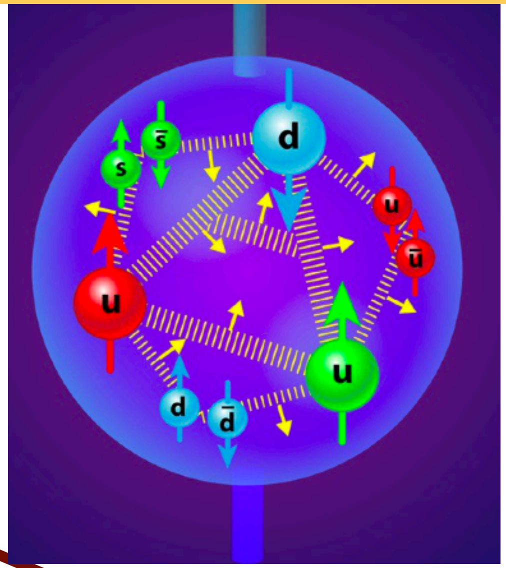

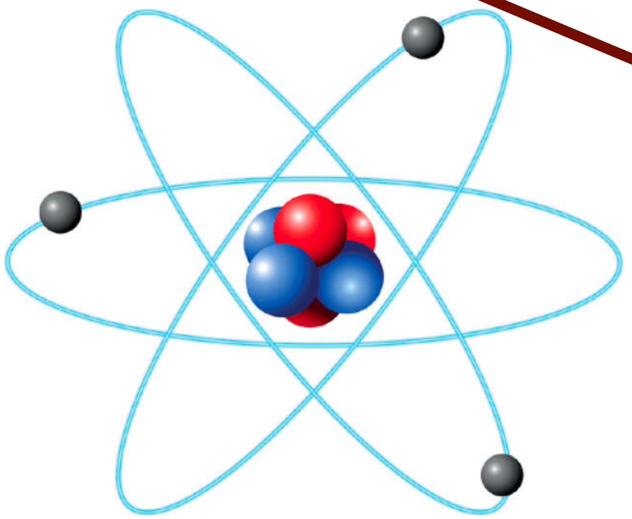

atomic nucleus (10^{-5} *L*)

Higgs boson production (10^{-8} *L*)

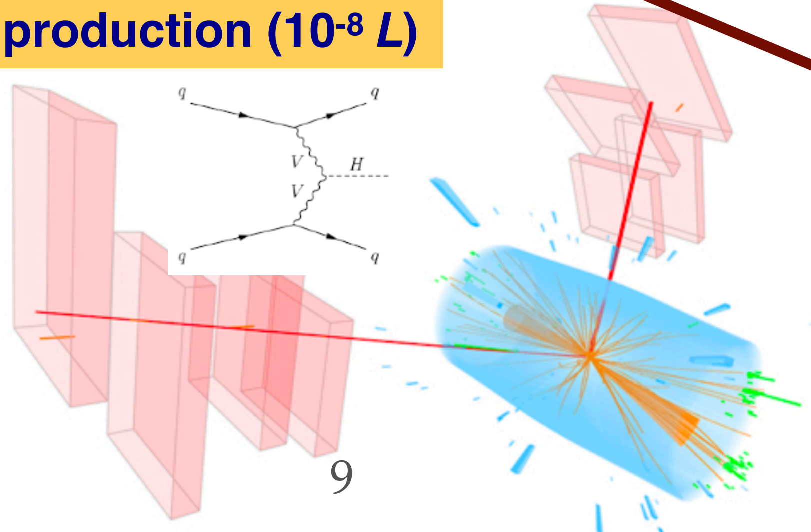

short distances
& large energies

Journeys into the Zeptospace

$$\text{lenght units} \rightarrow L = 10^{-10}\ \text{m}$$ (Bohr's radius, size of H atom)

$$\text{energy/mass units} \rightarrow E = 10^{9}\ \text{eV} = 1\ \text{GeV}$$ (mass of H atom, E=mc²)

long distances
& low energies

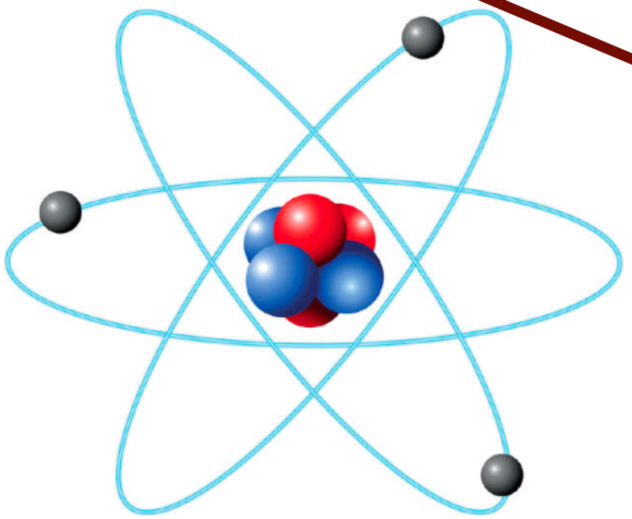

atomic nucleus
(10^{-5} *L*)

quark and gluon structure
of protons (10^{-6} *L*)

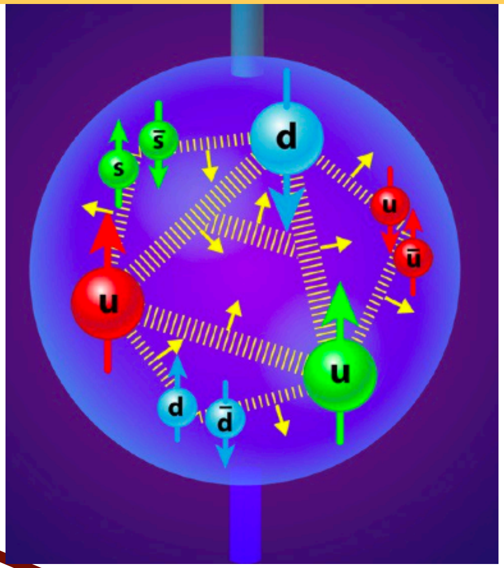

Higgs boson
production (10^{-8} *L*)

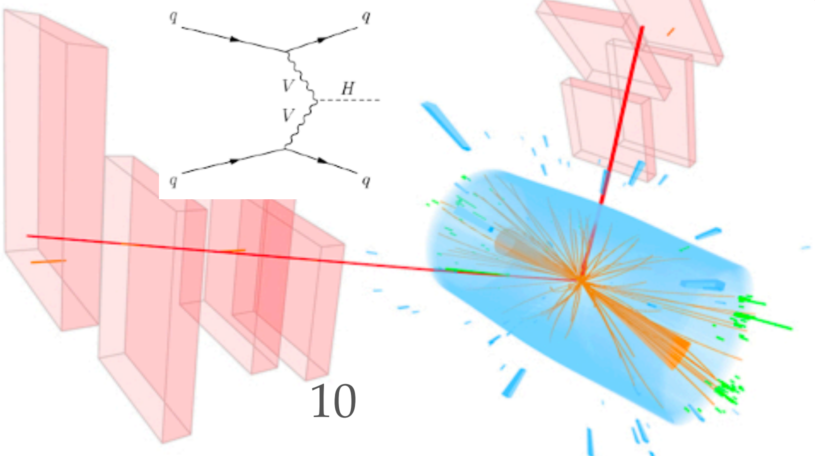

high-mass Drell-Yan (10^{-10} *L*)

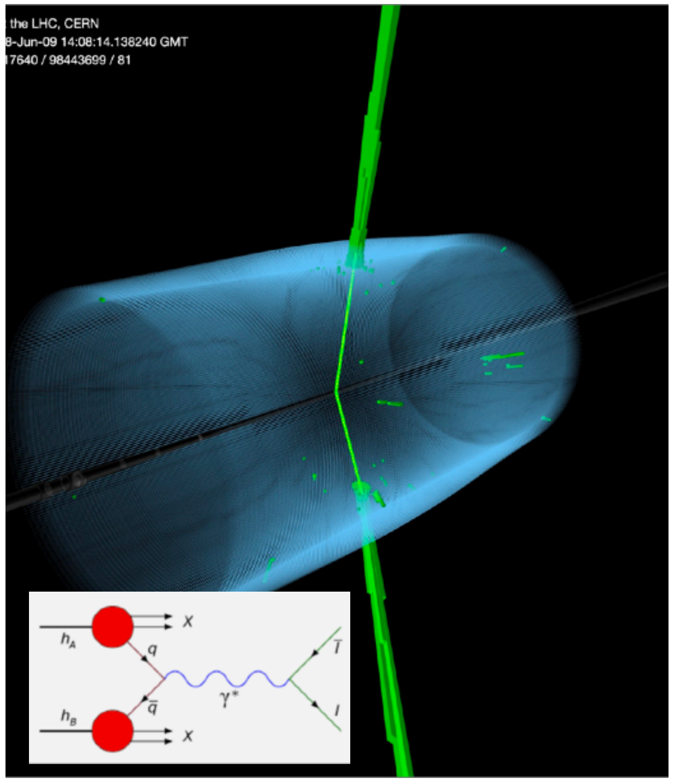

The Zeptospace!
10^{-11} L

short distances
& large energies

Journeys into the Zeptospace

$$\text{lenght units} \rightarrow L = 10^{-10}\ \text{m}$$ (Bohr's radius, size of H atom)

$$\text{energy/mass units} \rightarrow E = 10^{9}\ \text{eV} = 1\ \text{GeV}$$ (mass of H atom, $E=mc^2$)

long distances
& low energies

Rutherford's experiment

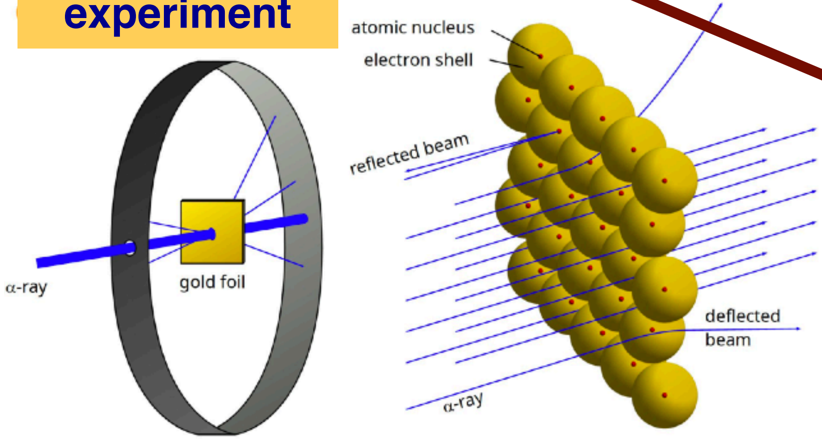

atomic nucleus (10^{-5} *L*, probe 10^{-3} E)

quark and gluon structure of protons (10^{-6} *L*, probe 10^{1} E)

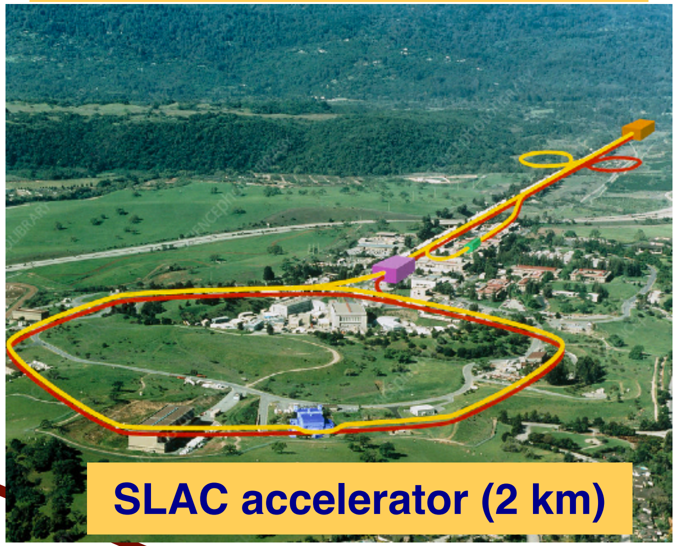

SLAC accelerator (2 km)

Higgs boson (10^{-8} *L*, probe 10^{3} *E*)

Drell-Yan (10^{-10} *L*, probe 10^{4} *E*)

Large Hadron Collider (27 km)

CMS experiment @ LHC

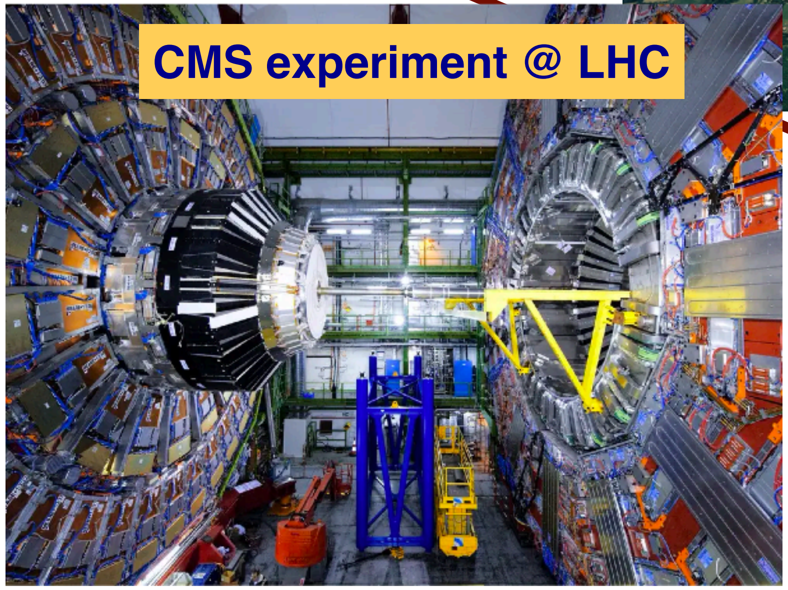

The Zeptospace!
10^{-11} L

short distances
& large energies

The Large Hadron Collider

The Electron-Ion Collider

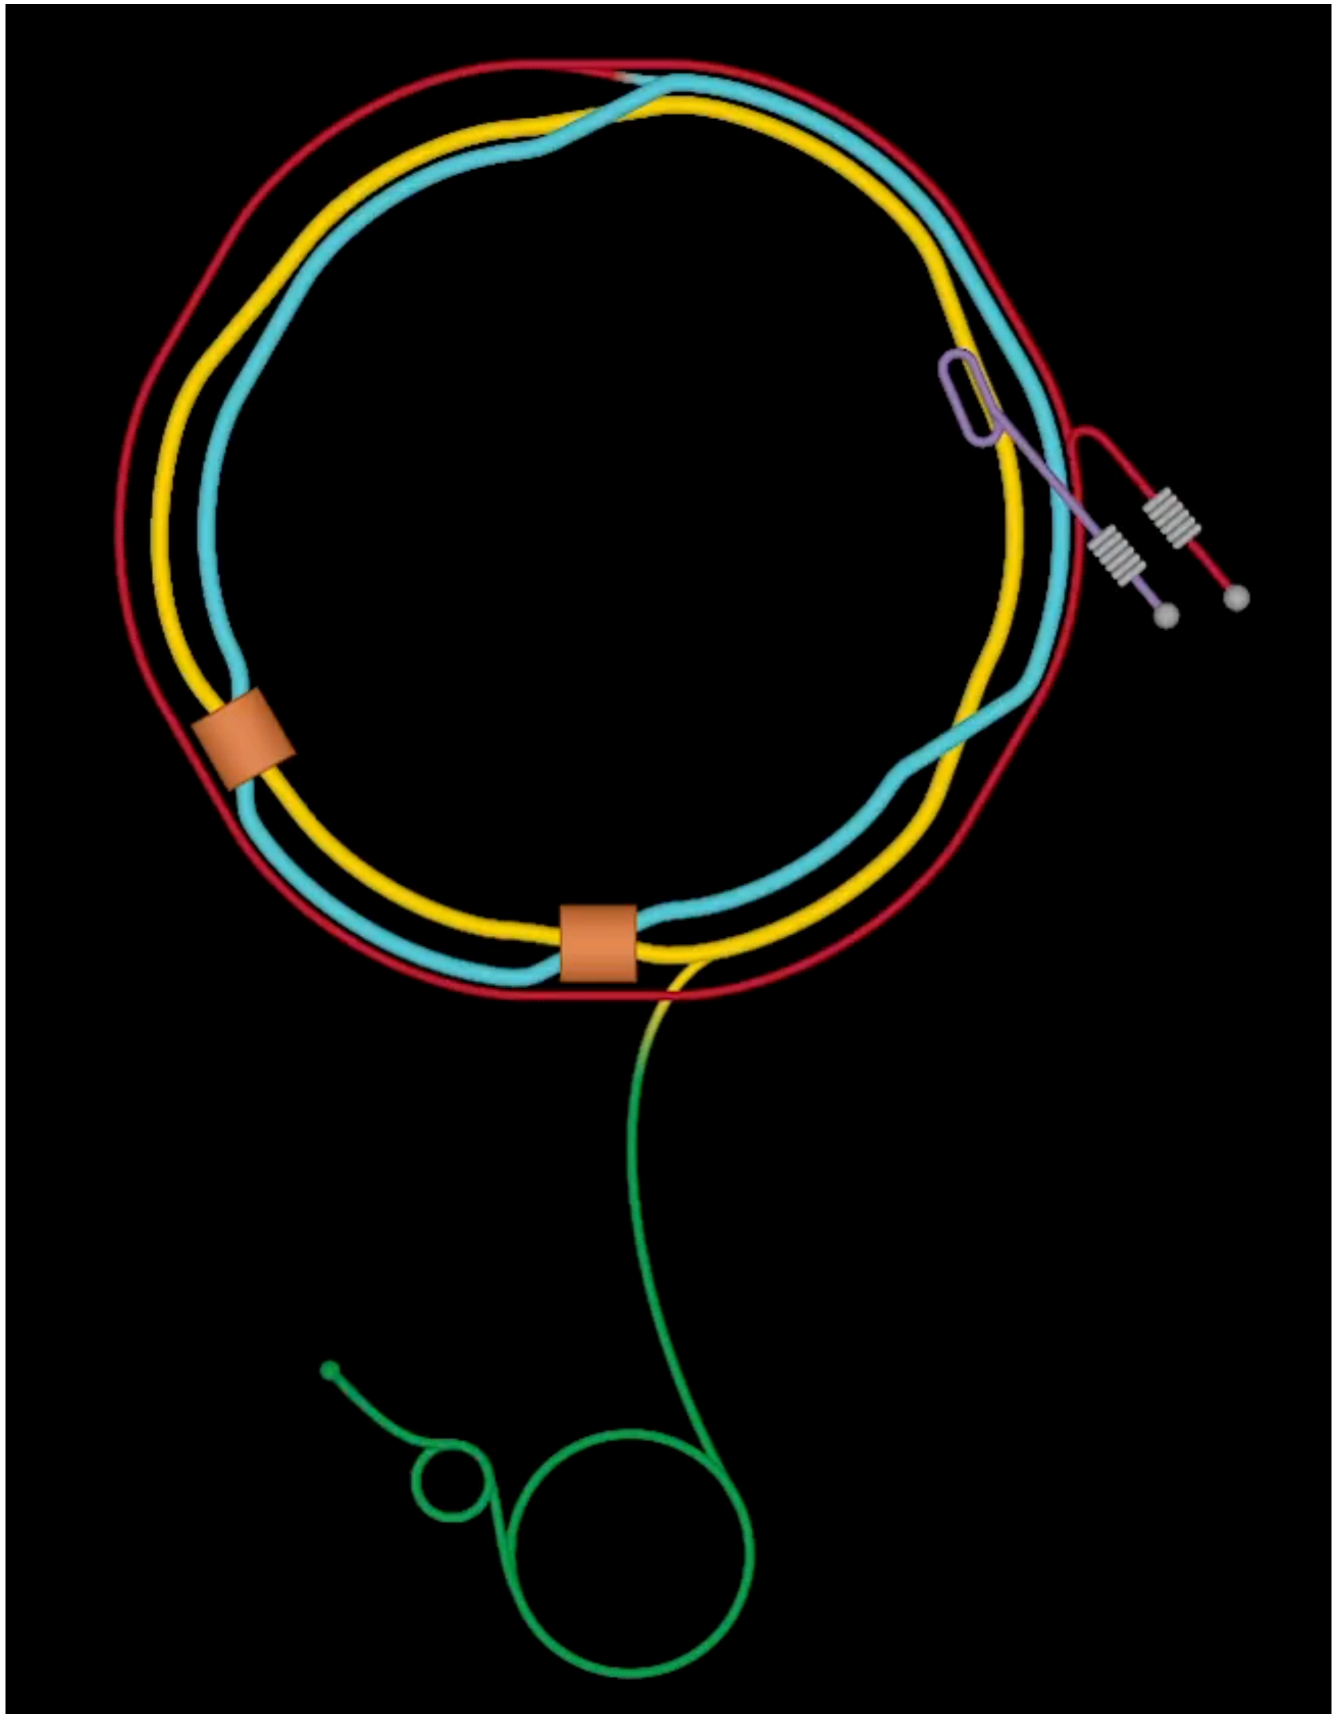

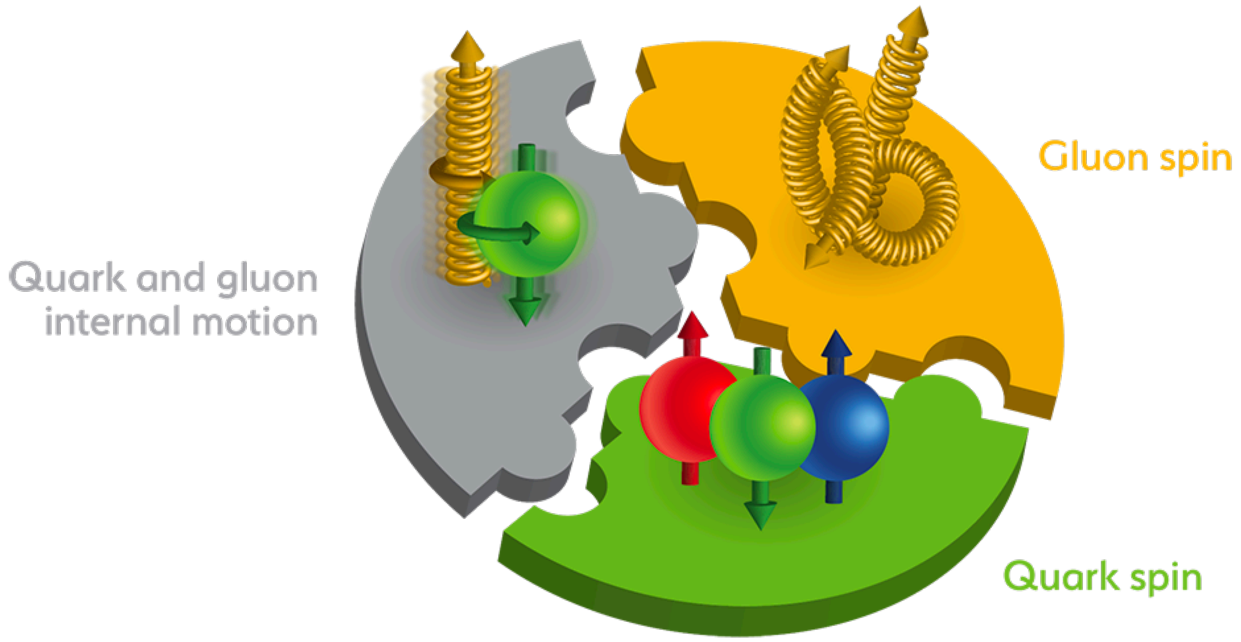

where does proton spin $s=1/2$ come from?

The Standard Model: a Success Story

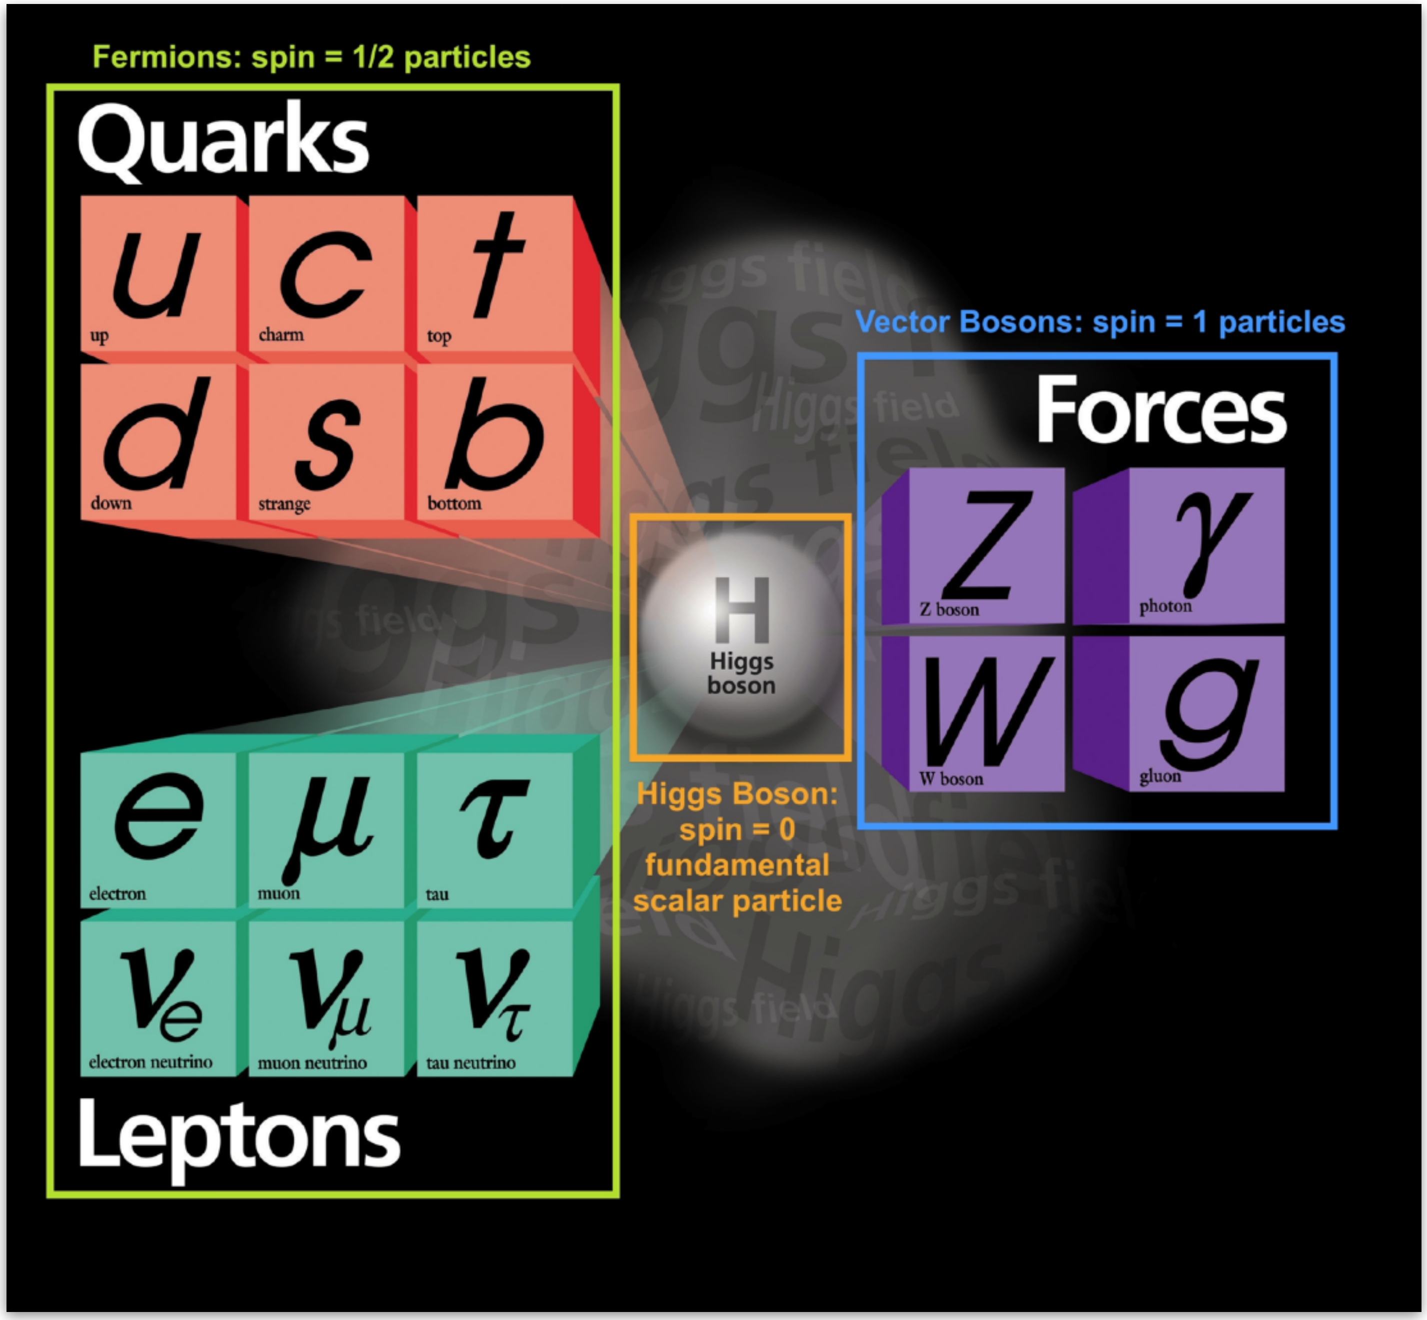

- **Robust** and **predictive** mathematical framework describing **elementary particles** and their (non-gravitational) interactions
- **Matter particles:** three families of quarks and leptons
- **Force mediators:** photon (QED), gluons (strong interaction), W & Z bosons (electroweak force)
- The Higgs field and its excitation, the Higgs boson: both a **completely new fundamental particle & interaction** (first non-gauge elementary force)

Do we really need all these particles and interactions?

The Standard Model: a Success Story

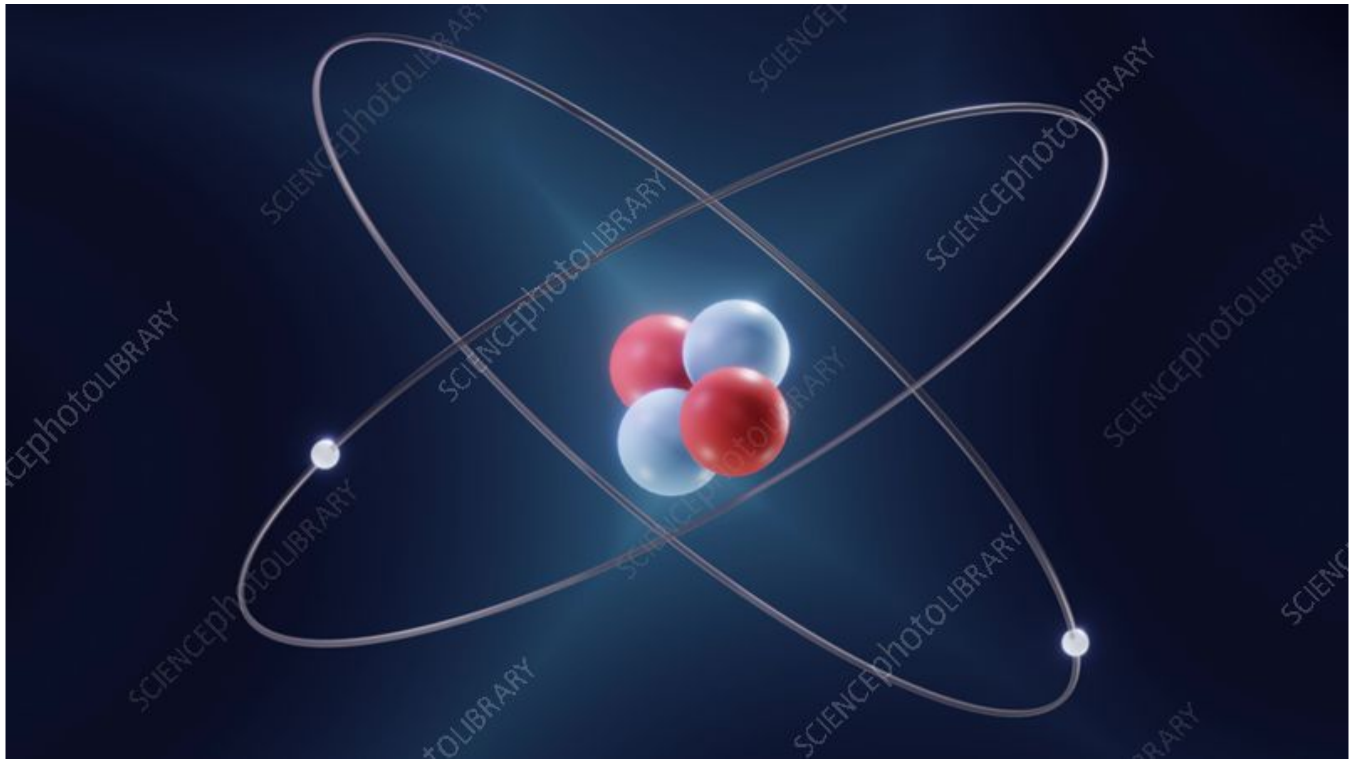

- **Everyday matter** composed exclusively of protons, electrons, and neutrons
- However, without the **full complexity of the SM**, the universe as we know it would be impossible!

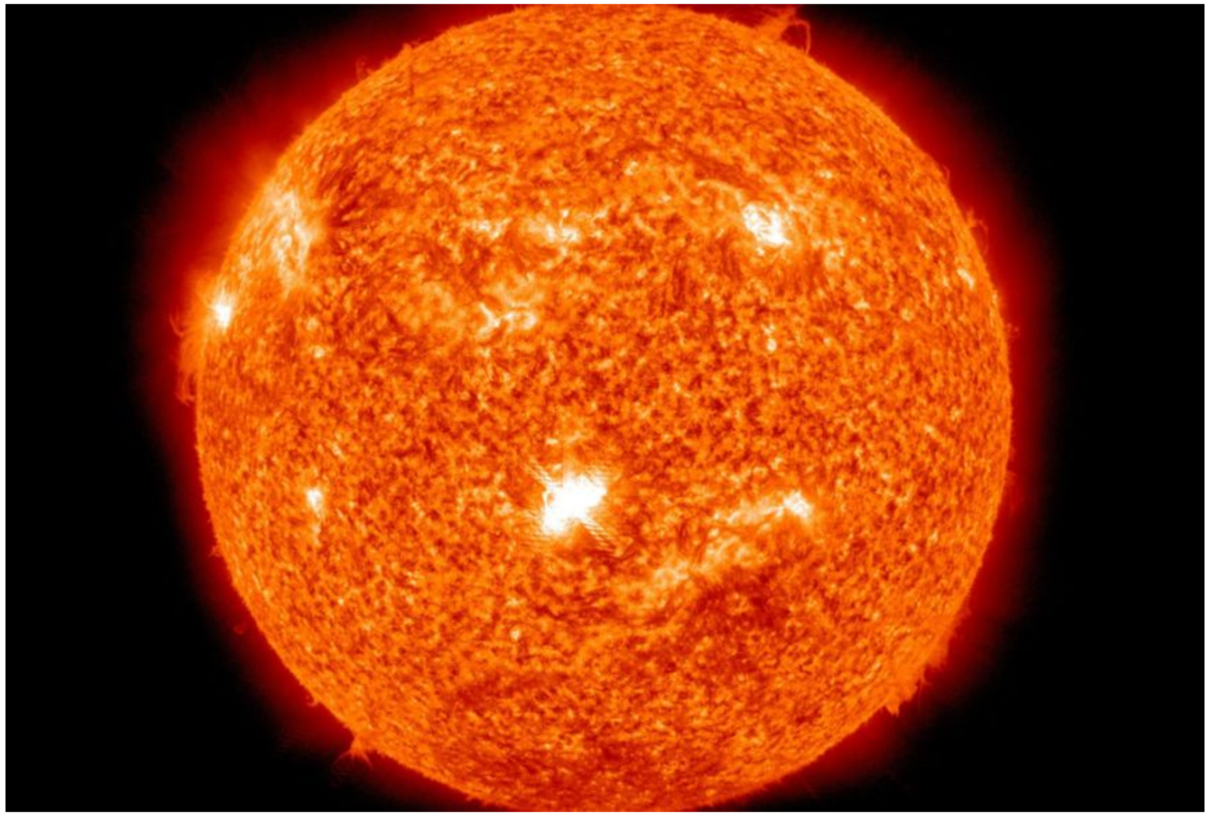

The Standard Model: a Success Story

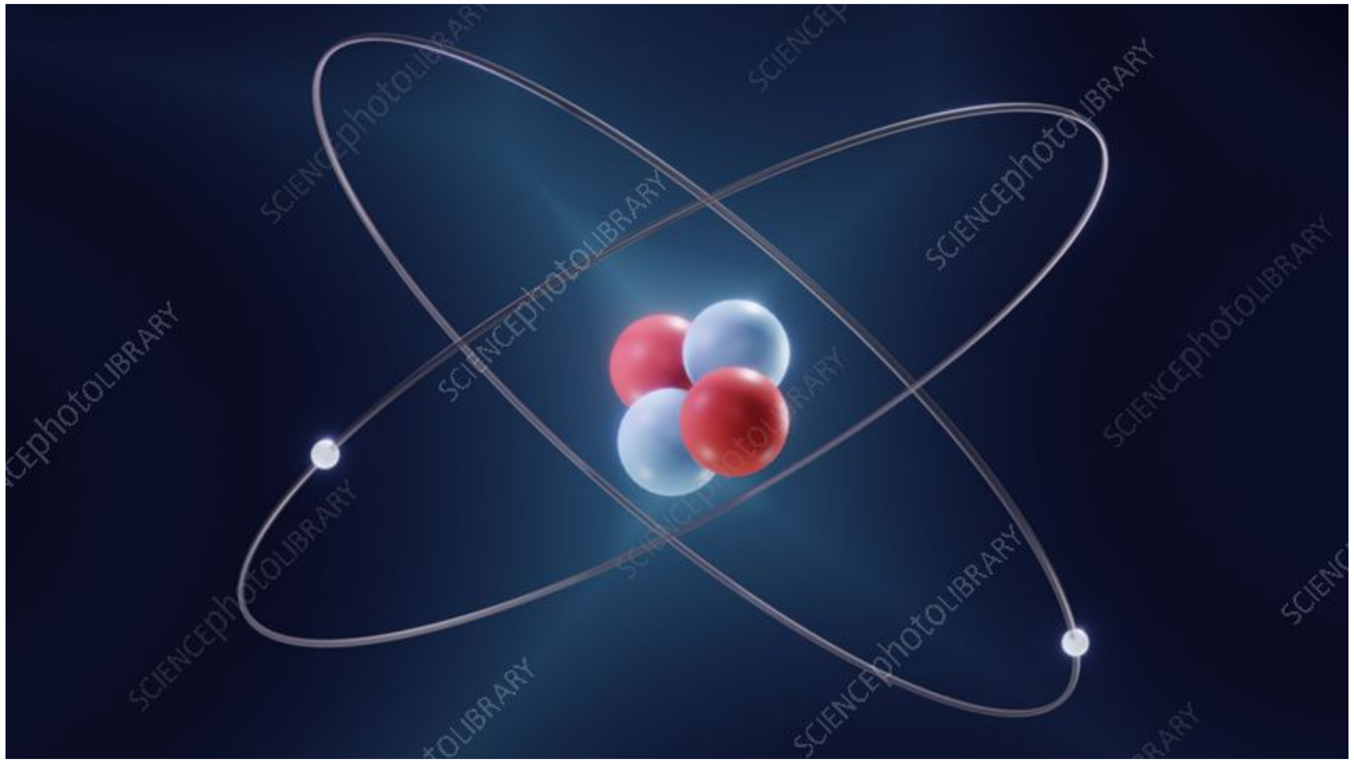

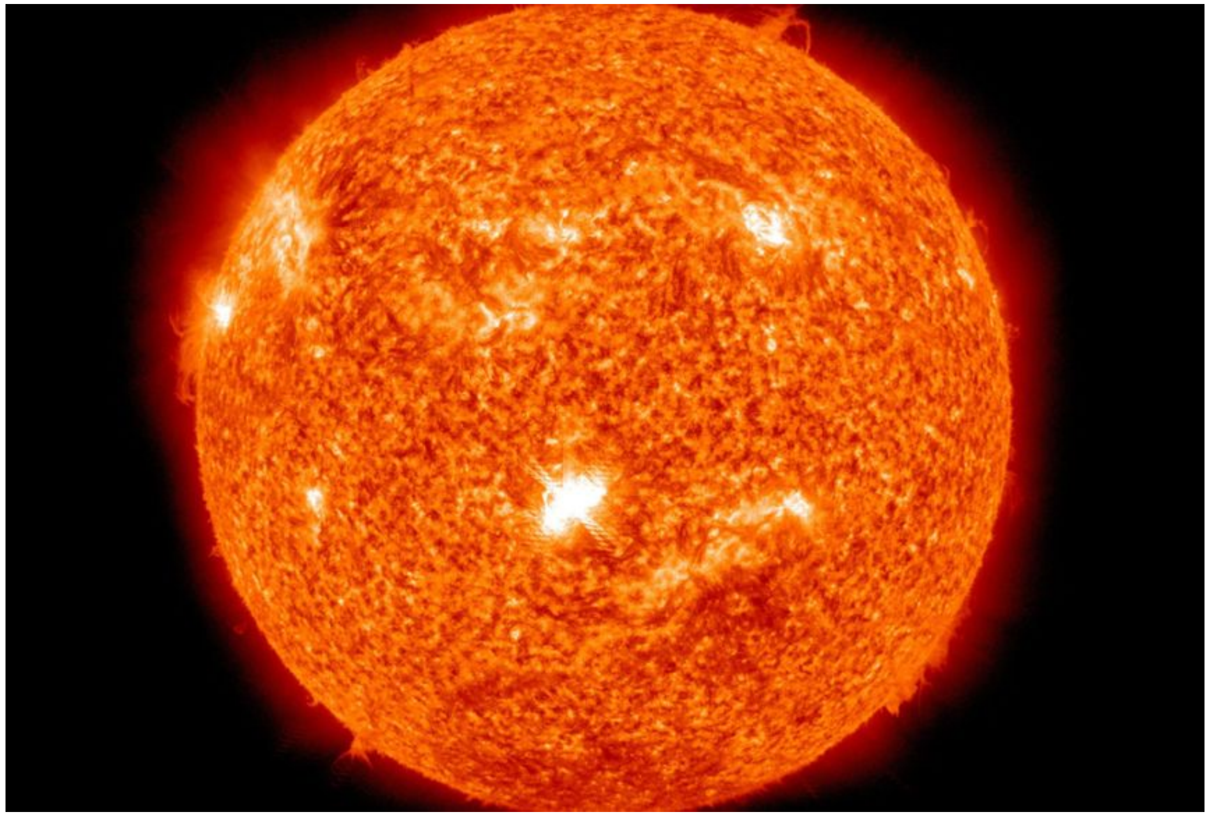

- **Everyday matter** composed exclusively of protons, electrons, and neutrons
- However, without the **full complexity of the SM**, the universe as we know it would be impossible!
- No nuclear fusion (and hence no stars!) without the **strong nuclear force**
- Matter would have fully annihilated with antimatter in early universe without CP violation, which **requires 3 quark generations**
- Without the Higgs boson, the **electron would be massless** and hence stable matter would not exist

The Standard Model: a Success Story

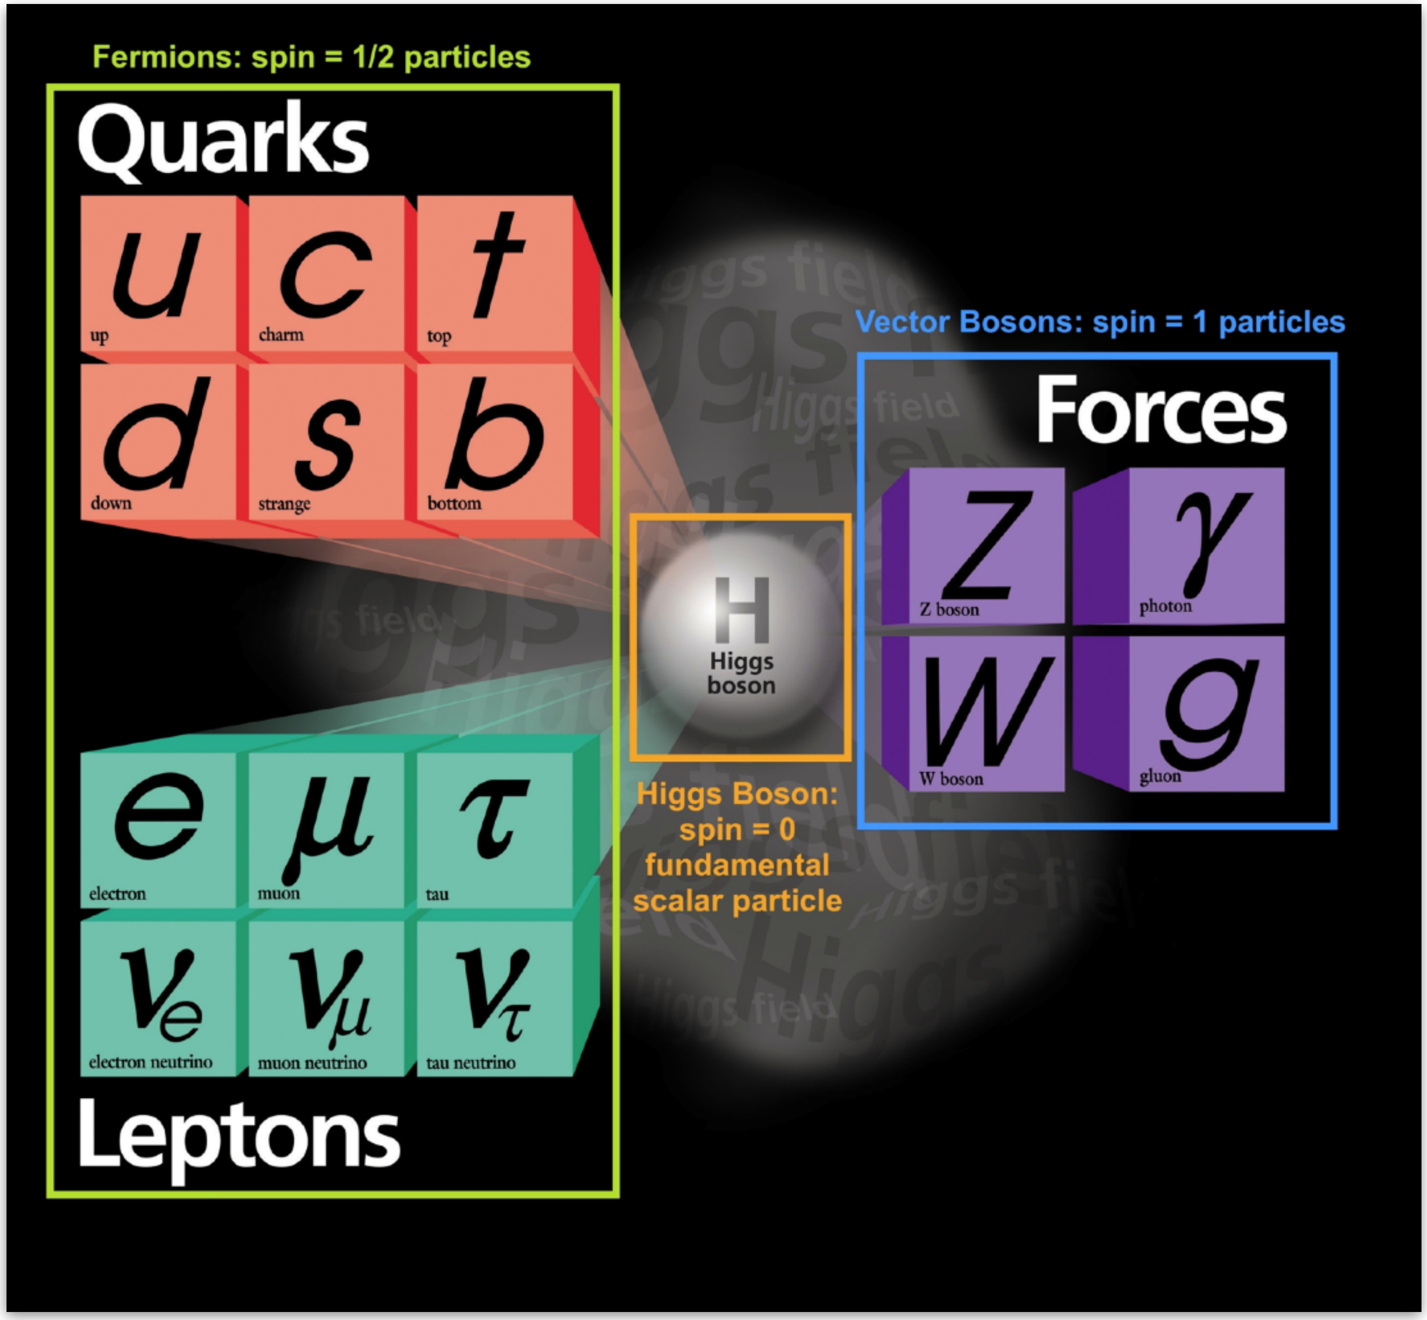

- **Robust** and **predictive** mathematical framework describing **elementary particles** and their (non-gravitational) interactions
- **Matter particles:** three families of quarks and leptons
- **Force mediators:** photon (QED), gluons (strong interaction), W & Z bosons (electroweak force)
- The Higgs field and its excitation, the Higgs boson: both a **completely new fundamental particle & interaction** (first non-gauge elementary force!)

Describes an incredibly wealth of measurements with astonishing precision: a major triumph of modern science

The Standard Model: not the Full Story

Is the Higgs boson **elementary or composite**? Do other fundamental scalars exist?

Is there a more fundamental description of the **EWSB mechanism**? What sets strength of Yukawa interactions?

What determines the **pattern of masses and couplings** of the SM particles?

$$m_t/m_e \sim 4 \times 10^5$$

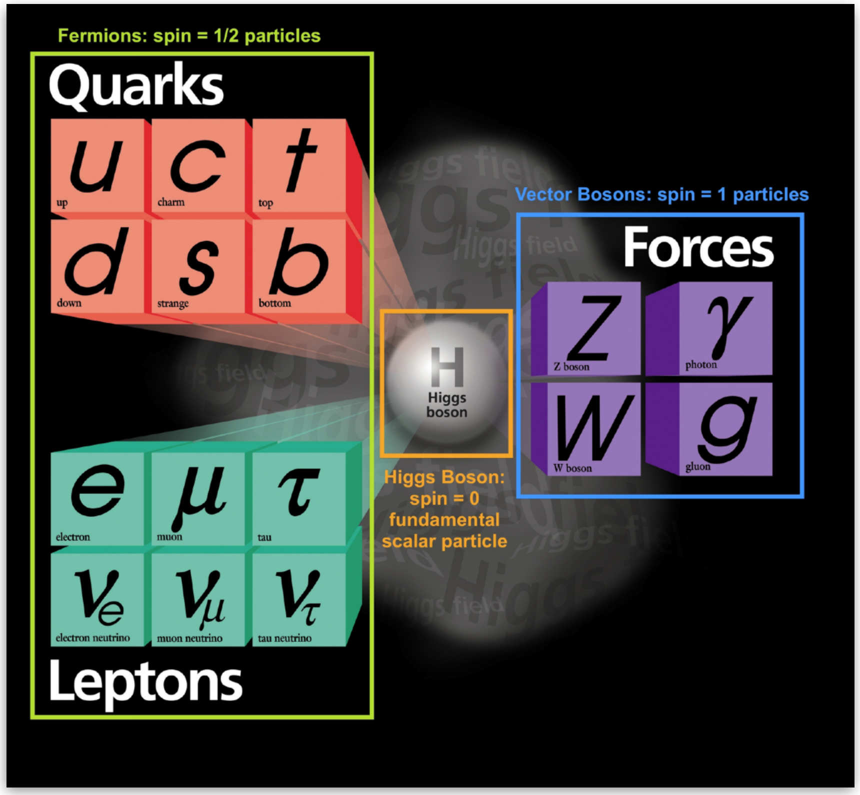

Does the Higgs mechanism provide a mass for all **1st and 2nd generation fermions**? And for **neutrinos**?

The Standard Model: not the Full Story

Is the Higgs boson **elementary or composite**? Do other fundamental scalars exist?

Is there a more fundamental description of the **EWSB mechanism**? What sets strength of Yukawa interactions?

What determines the **pattern of masses and couplings** of the SM particles?

Which mechanisms explains the observed **matter/antimatter asymmetry** of the universe?

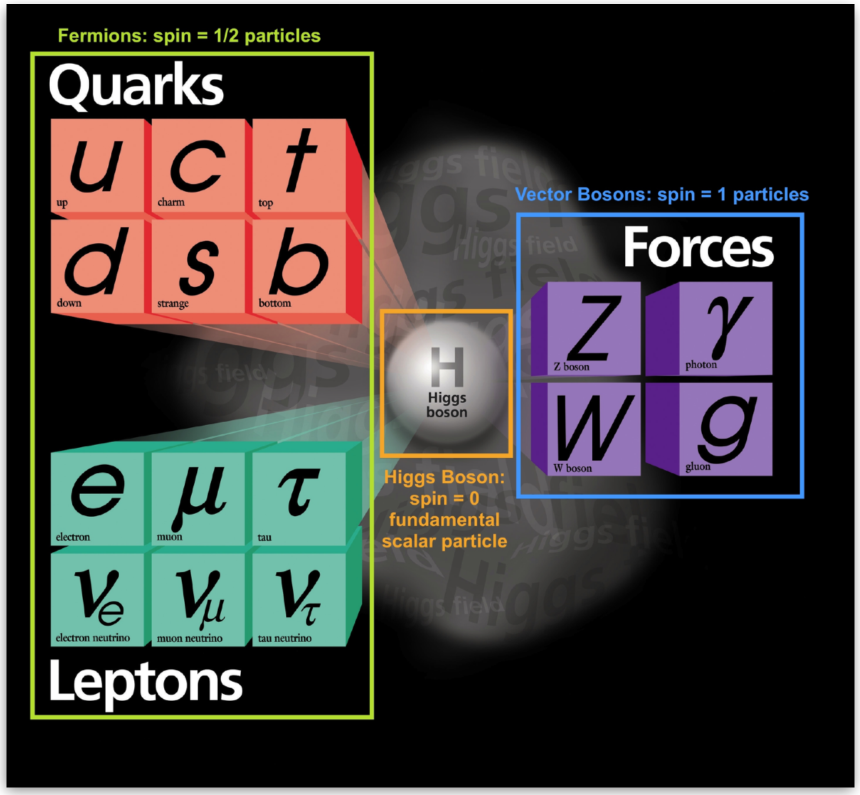

Does the Higgs mechanism provide a mass for all **1st and 2nd generation fermions**? And for **neutrinos**?

Is there a **particle-based description of Dark Matter**? Do Dark Sectors exist?

What is the correct **quantum mechanical description of gravity**? Can we test it in the lab or the cosmos?

and many others + new particles/forces may exist without serving any obvious ``purpose"

Innumerable extensions of the SM have been proposed. Despite extensive efforts, no clear pathway to physics Beyond the SM has been identified

Fundamental Physics: a thriving field

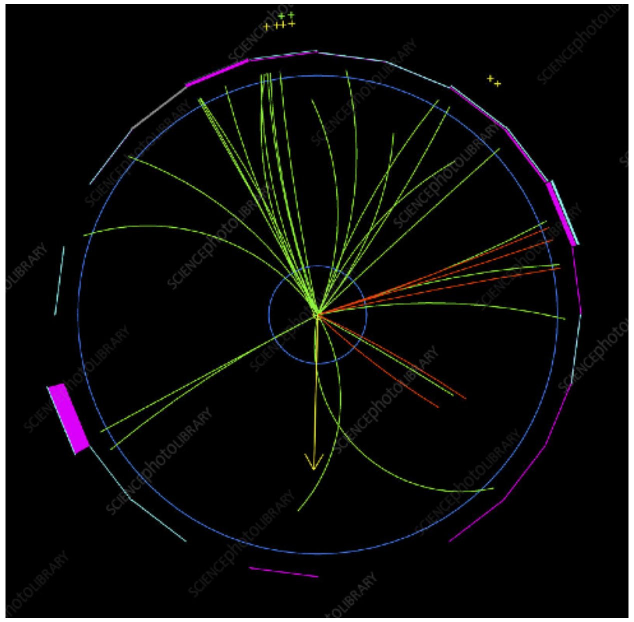

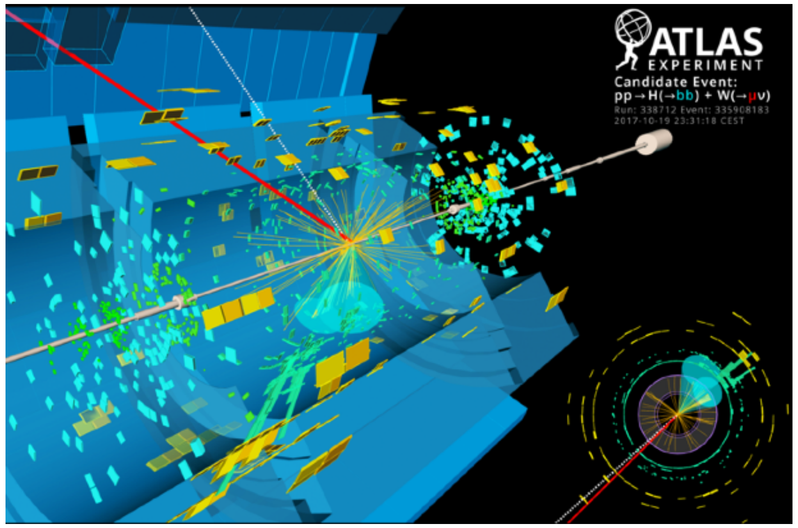

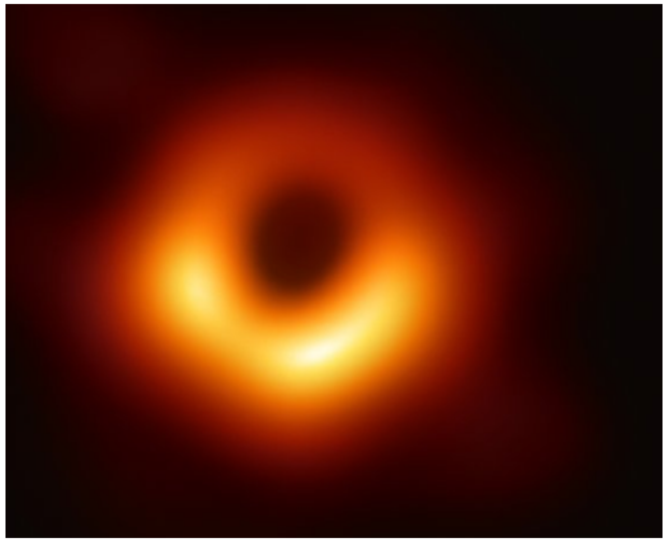

1994: top quark discovery

1998: discovery of neutrino masses

2012: Higgs boson discovery

2019: first Black Hole direct observation

1998: discovery of cosmic acceleration

2015: discovery of gravitational waves

2030+: discovery of Higgs self-interactions?

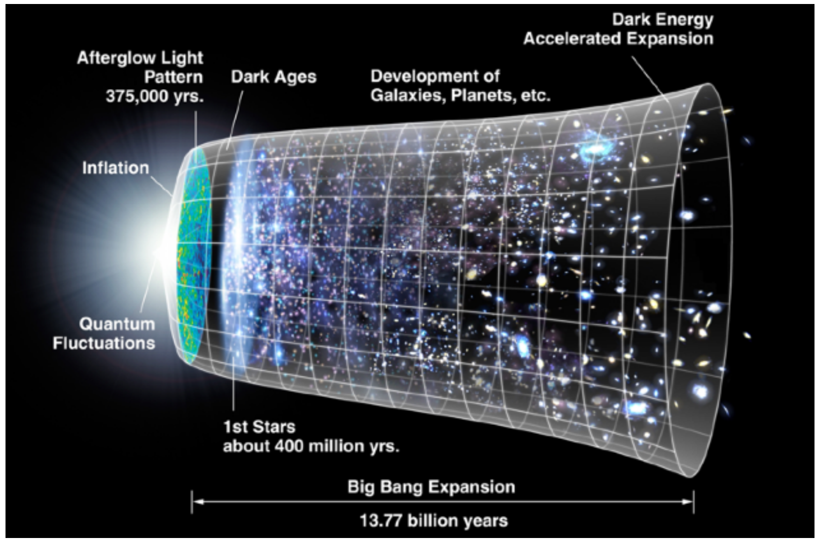

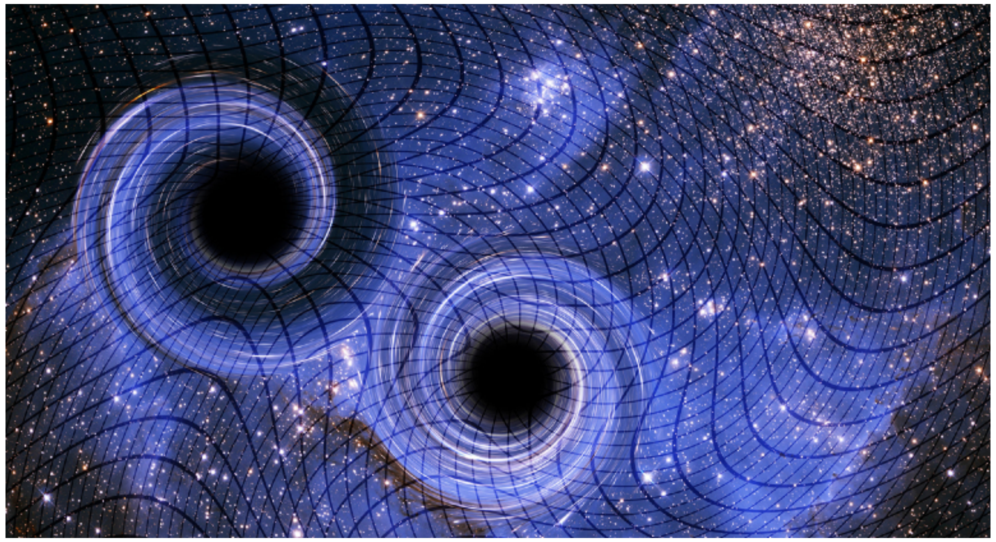

ML & AI for HEP


THIS IS YOUR MACHINE LEARNING SYSTEM?
YUP! YOU POUR THE DATA INTO THIS BIG PILE OF LINEAR ALGEBRA, THEN COLLECT THE ANSWERS ON THE OTHER SIDE.
WHAT IF THE ANSWERS ARE WRONG?
JUST STIR THE PILE UNTIL THEY START LOOKING RIGHT.
DATA
ANSWERS

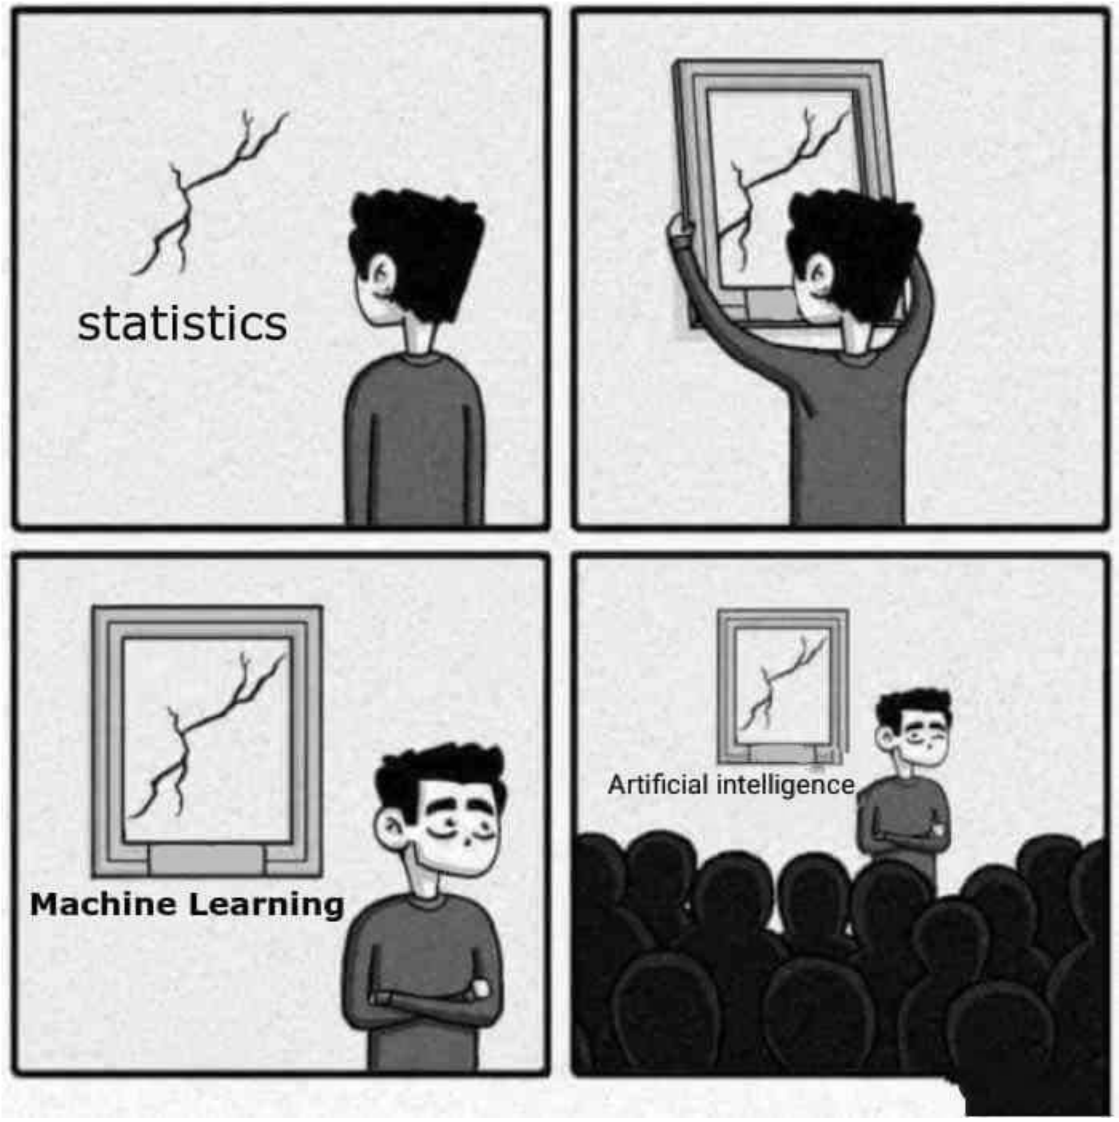
statistics
Machine Learning
Artificial intelligence

ML & AI for HEP

- Machine Learning has been used in HEP applications **since the early 90s**
- **Classification:** learn **decision boundaries** to classify events into different categories
- **Regression:** learn **physical behaviour from data** without model assumptions
- **Inference:** learn **forward models** from simulations for parameter inference

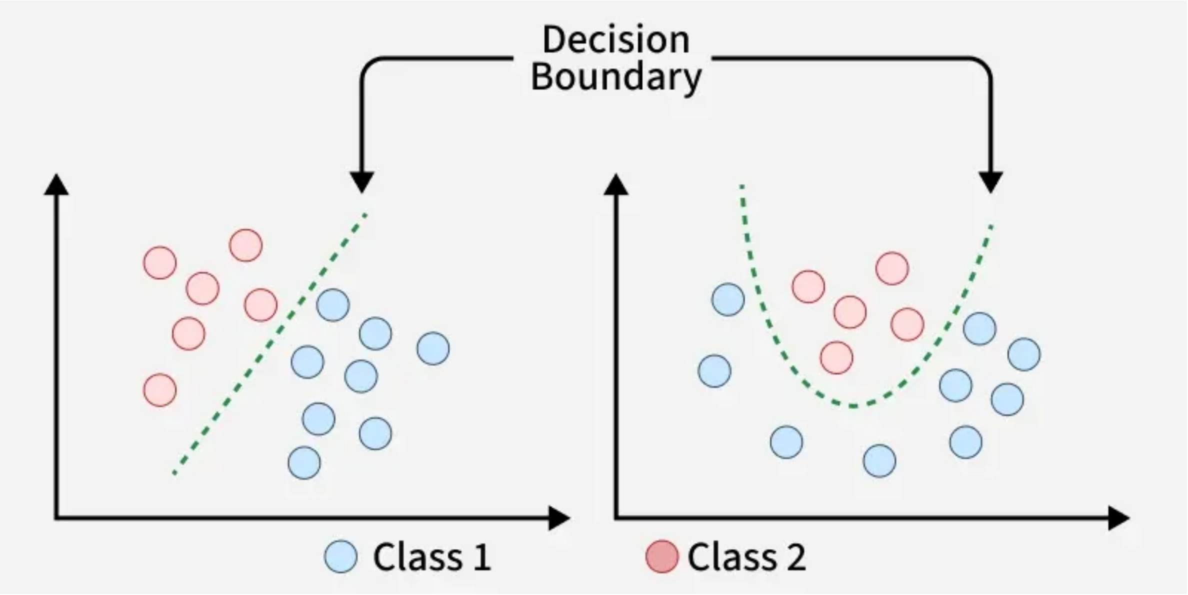

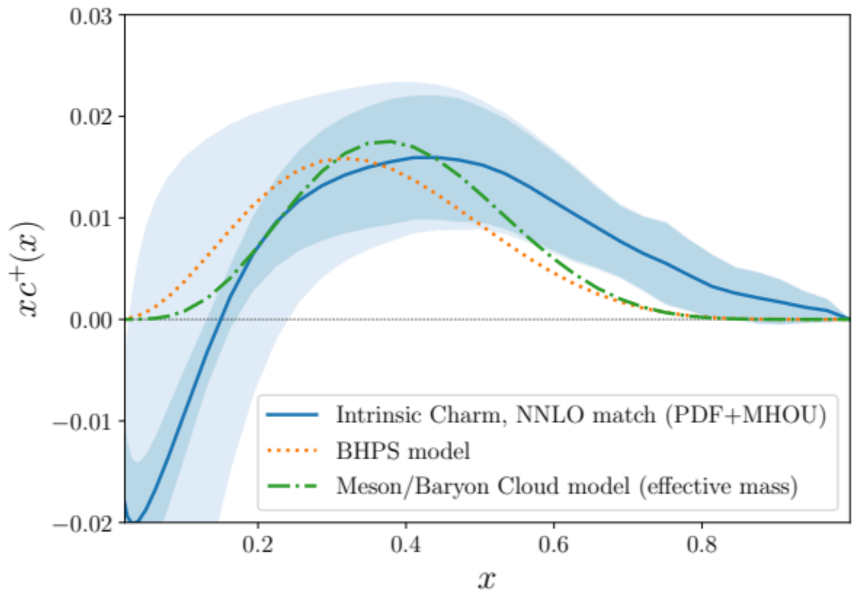

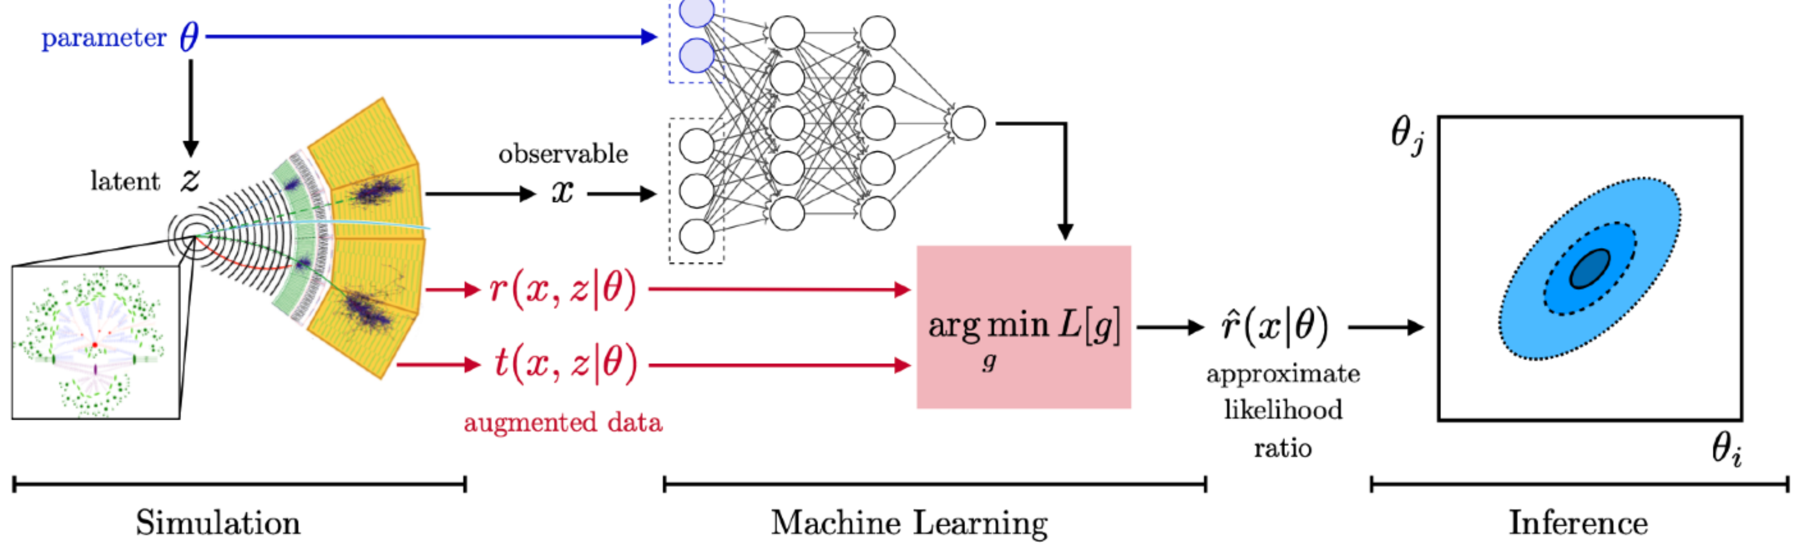

ML & AI for HEP

- **ML for regression**: discovery of intrinsic charm quarks in the proton
- **ML for inference**: optimally sensitive observables for proton structure
- **ML for classification**: towards evidence for Higgs self-interactions
- The next frontier: **Agentic AI for HEP**

Machine Learning and Proton Structure

The hidden life of protons

the proton is a beautiful example of the richness of quantum mechanics:
what a **proton is** depends on the **resolution with which we examine it**!

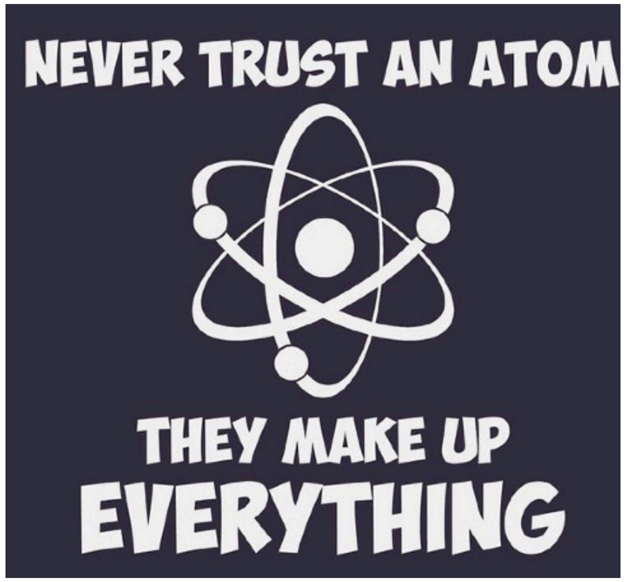

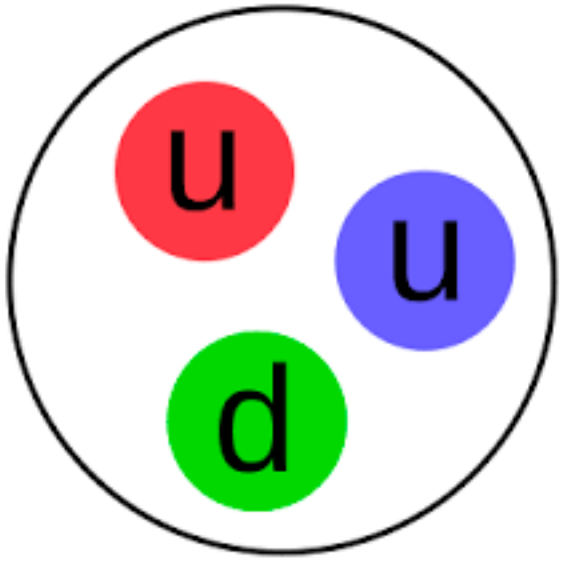

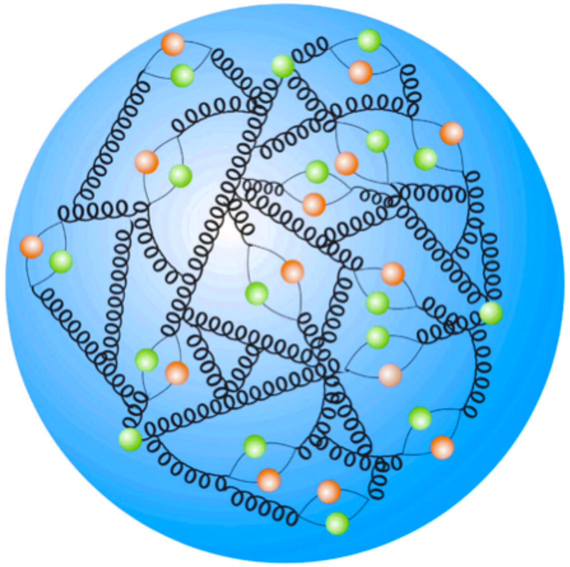

long distances / low energies → *short distances / high energies*

a point particle	3 valence quarks	sea quarks, gluons	heavy quarks, photons, leptons, Higgs bosons
$E \ll 1$ GeV	$E \sim$ few GeV	$E \sim 50$ GeV	$E \sim 1$ TeV

The hidden life of protons

credit: *visualising the proton*, Arts at MIT (https://arts.mit.edu/visualizing-the-proton/)

x: fraction of energy carried by proton constituents: the quarks and gluons

The hidden life of protons

The “charm quark” is a “heavier cousin” of the well-known up quark: same charge, etc., but **heavier than the proton itself**

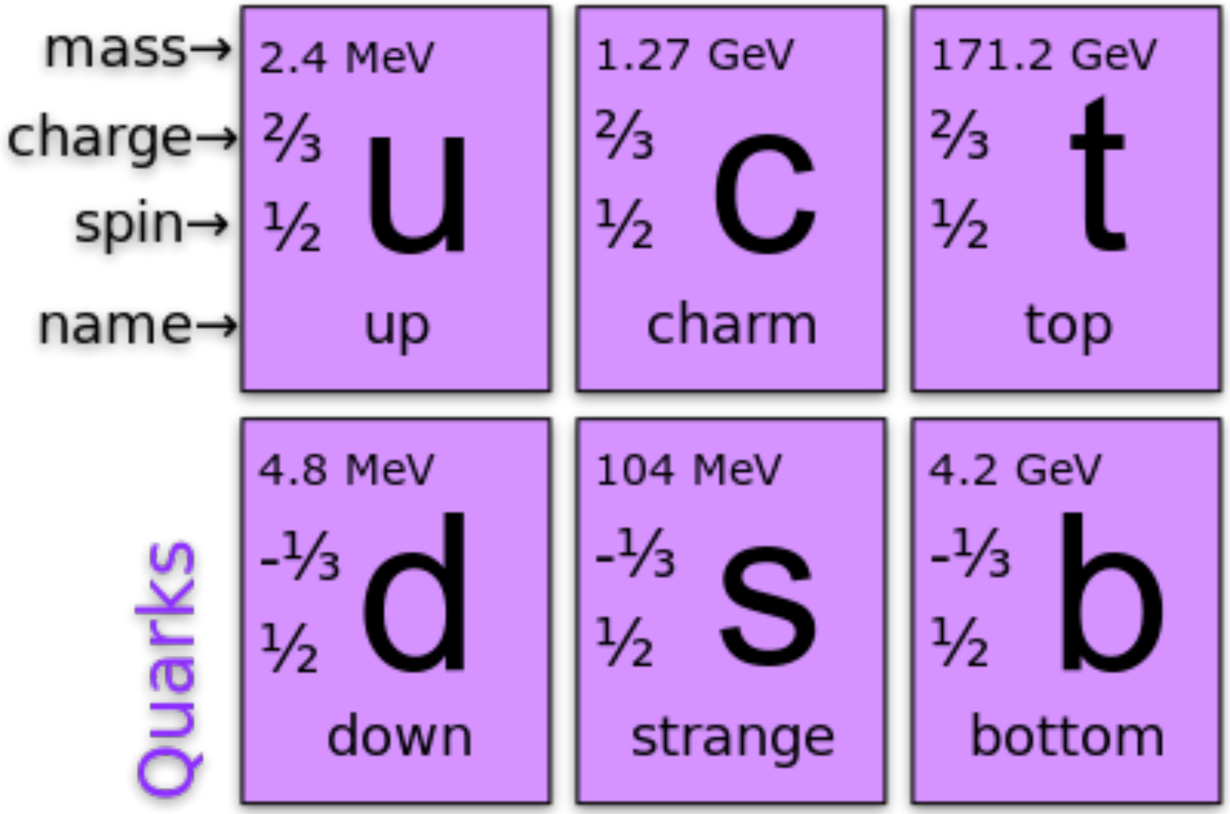

Quarks

	u	c	t
mass→	2.4 MeV	1.27 GeV	171.2 GeV
charge→	⅔	⅔	⅔
spin→	½	½	½
name→	up	charm	top

	d	s	b
mass	4.8 MeV	104 MeV	4.2 GeV
charge	-⅓	-⅓	-⅓
spin	½	½	½
name	down	strange	bottom

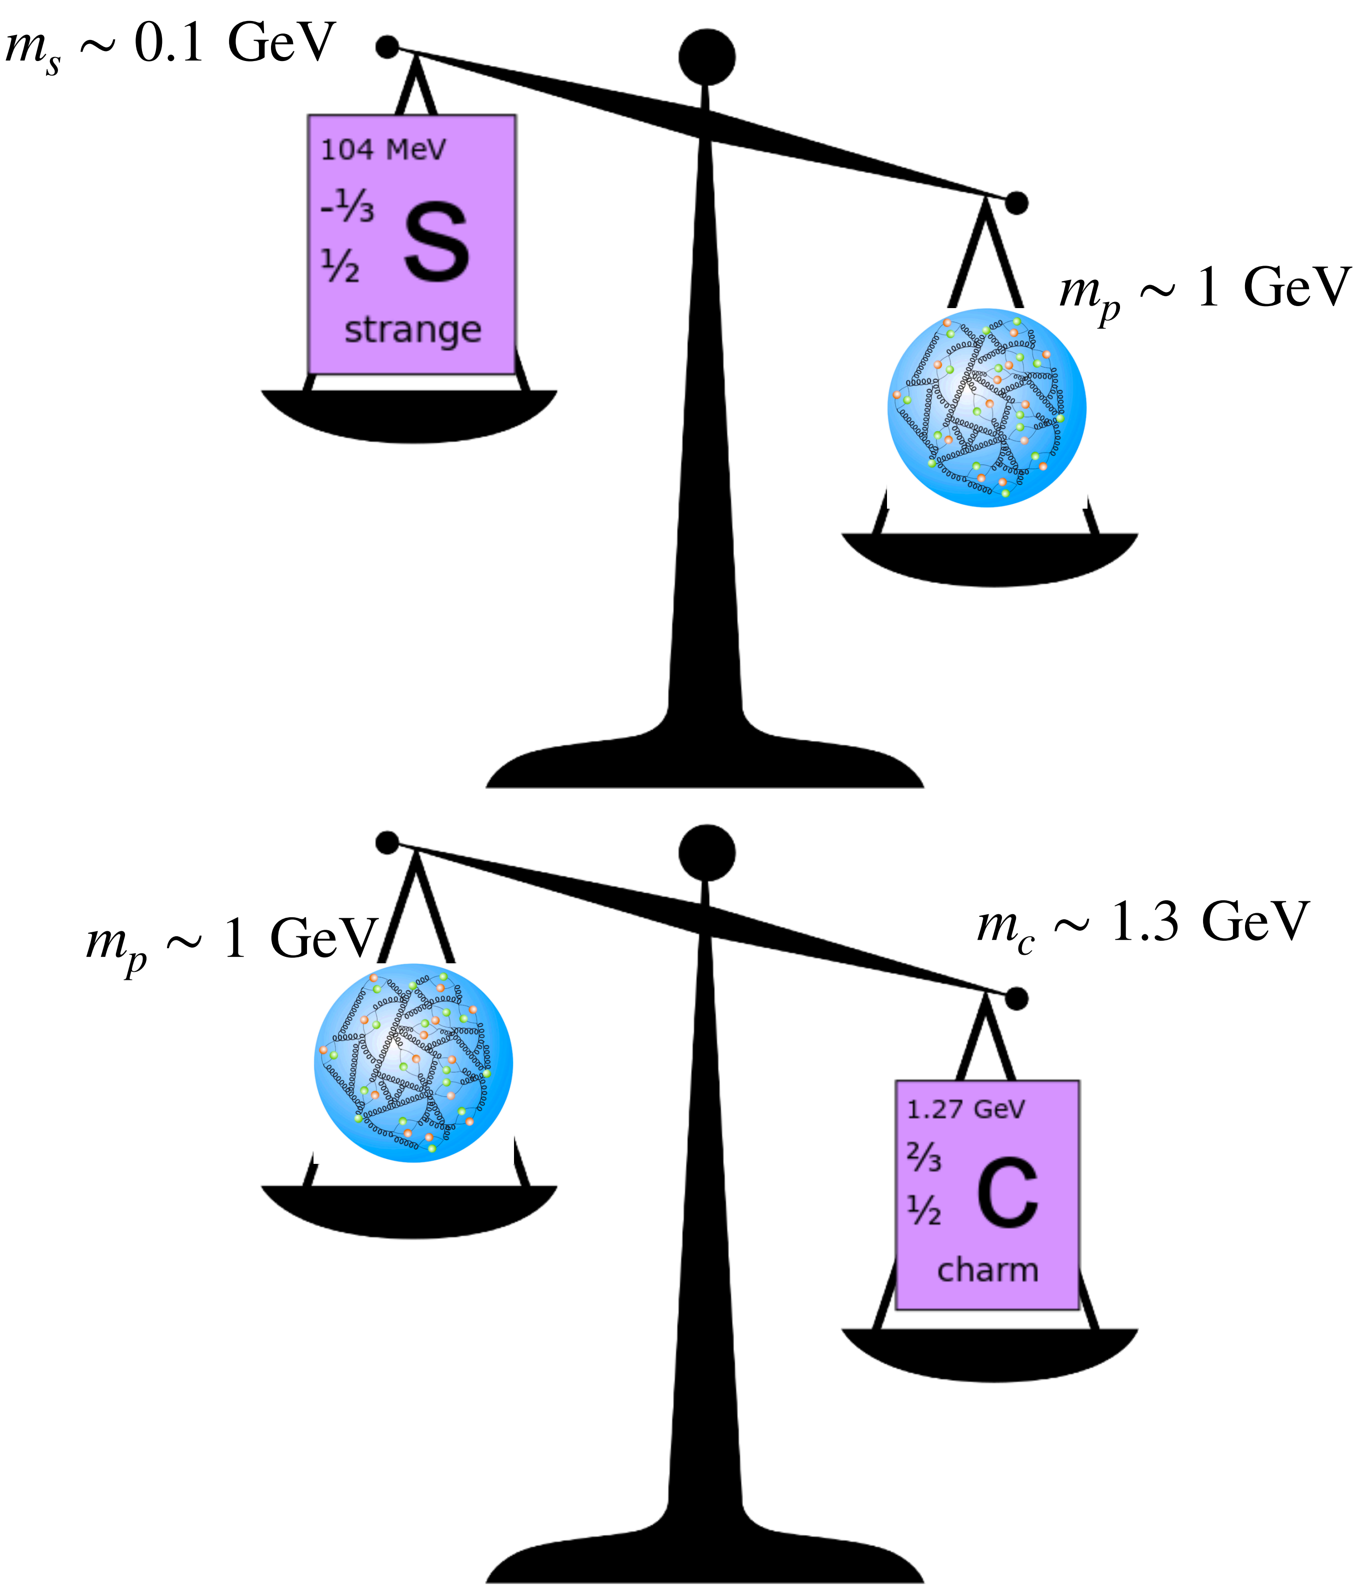

Would it be crazy to think that protons contain charm quarks?

What does Quantum Mechanics say?

The hidden life of protons

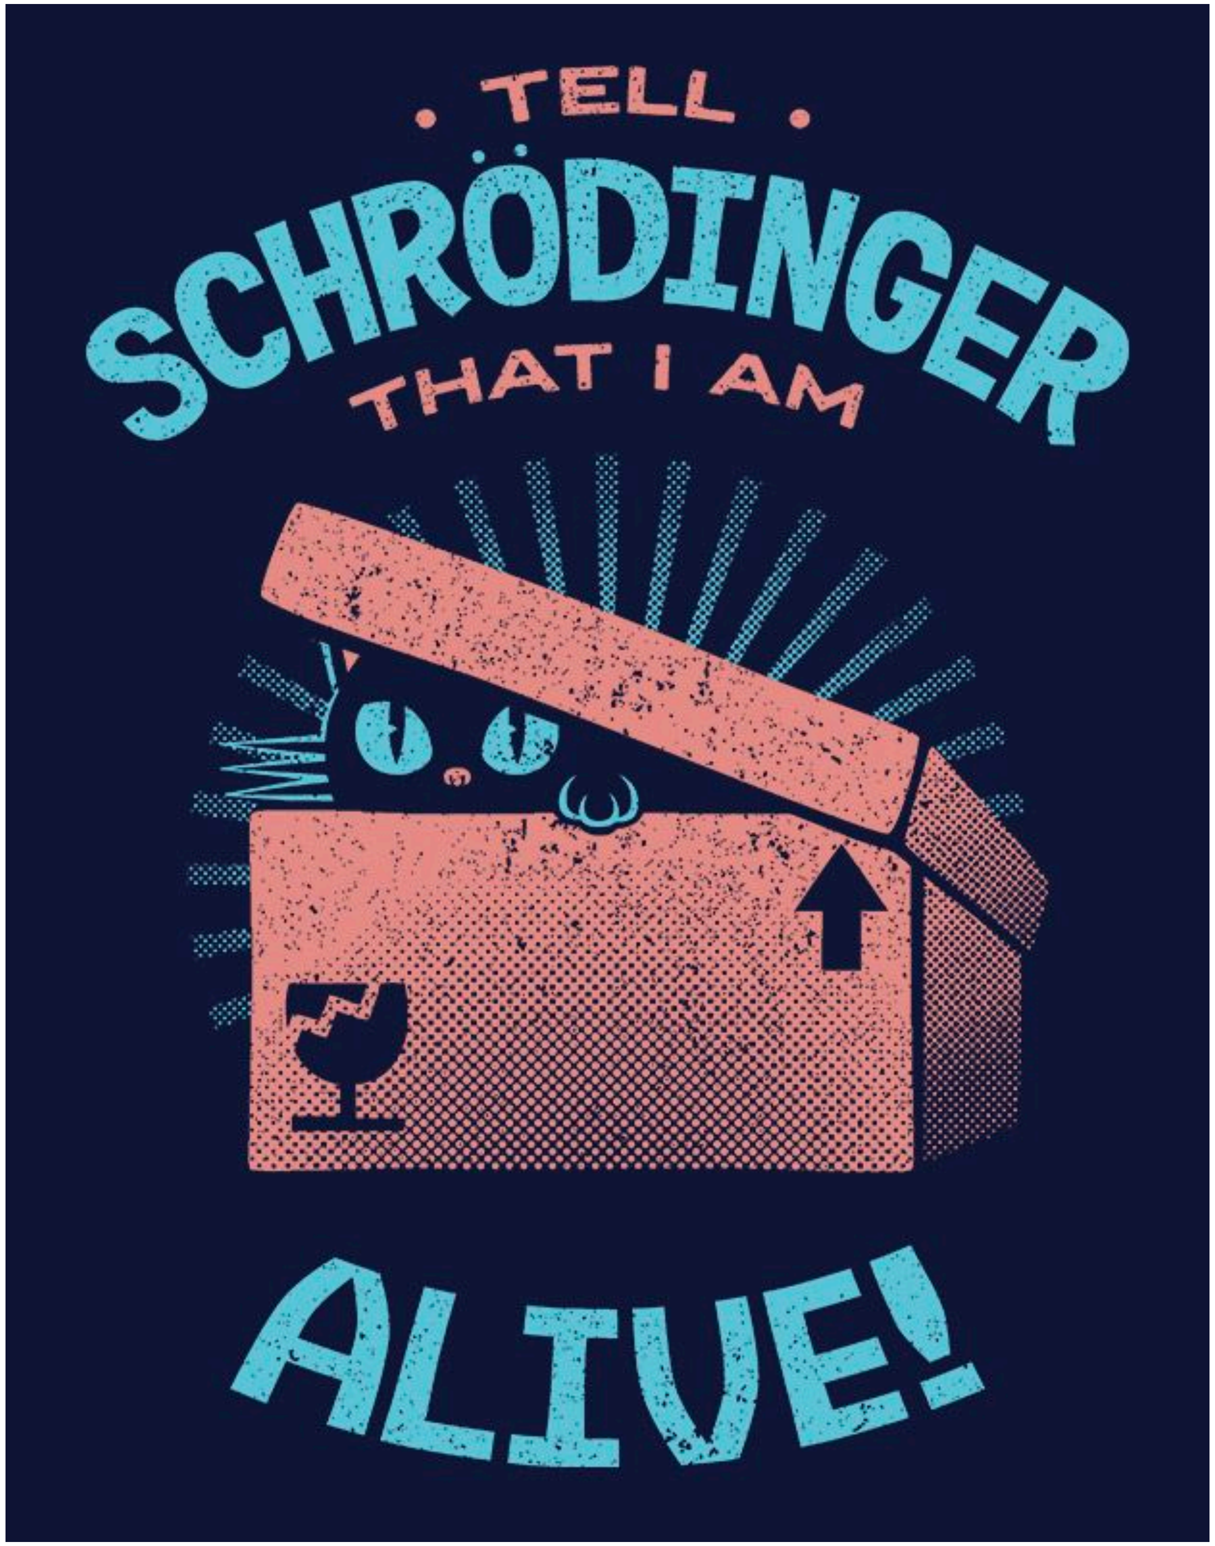

Textbook QM: quantum superposition

$$\big|\mathrm{Cat}\rangle \approx \big|\mathrm{Alive}\rangle + \big|\mathrm{Dead}\rangle$$

The same holds for proton?

The hidden life of protons

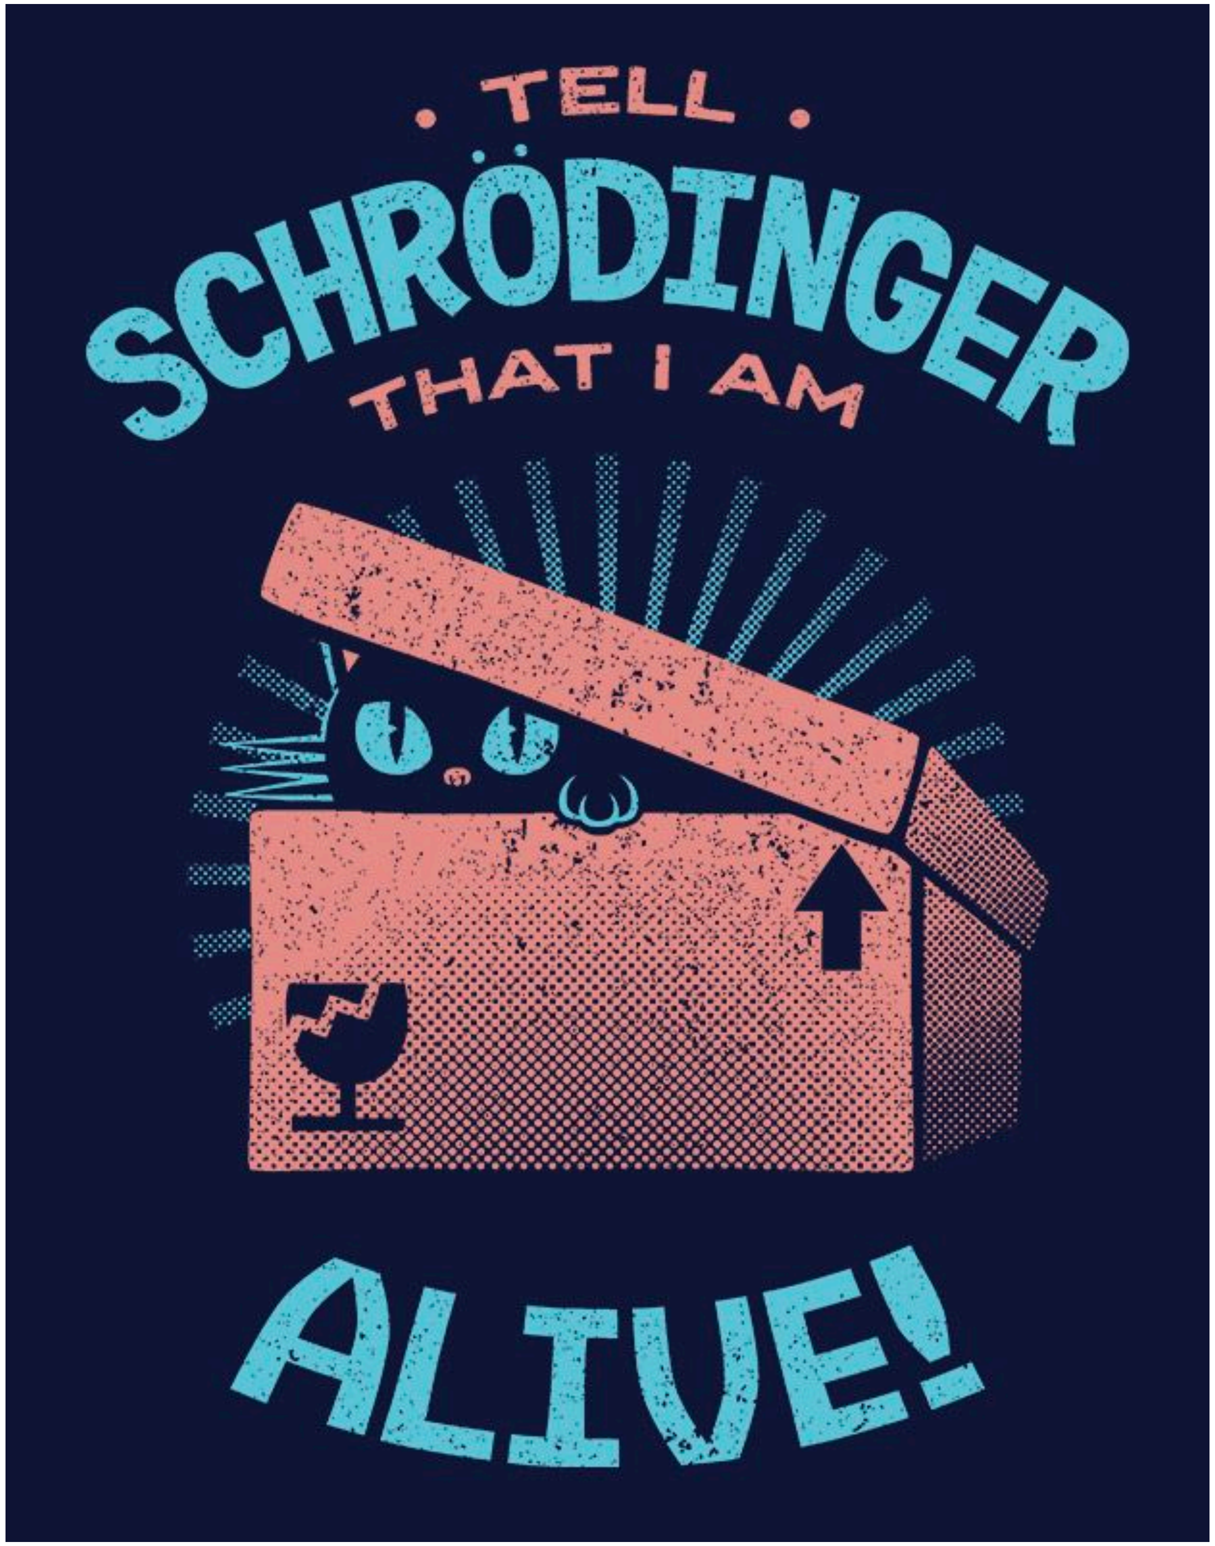

Textbook QM: quantum superposition

$$|\text{Cat}\rangle \approx |\text{Alive}\rangle + |\text{Dead}\rangle$$

The same holds for proton: it can, in principle, host constituents heavier than itself but **only for a short amount of time**

??

$$|\text{Proton}\rangle \approx 0.995\,|\text{uud}\rangle + 0.1\,|\text{uudc}\bar{\text{c}}\rangle$$

Light, high chance

``Regular” proton

Very heavy, small chance

Rare quantum fluctuation

$$\langle m_p\rangle = 0.99 m_{uud} + 0.01 m_{uudc\bar{c}}$$

credits: https://www.quantamagazine.org/inside-the-proton-the-most-complicated-thing-imaginable-20221019/

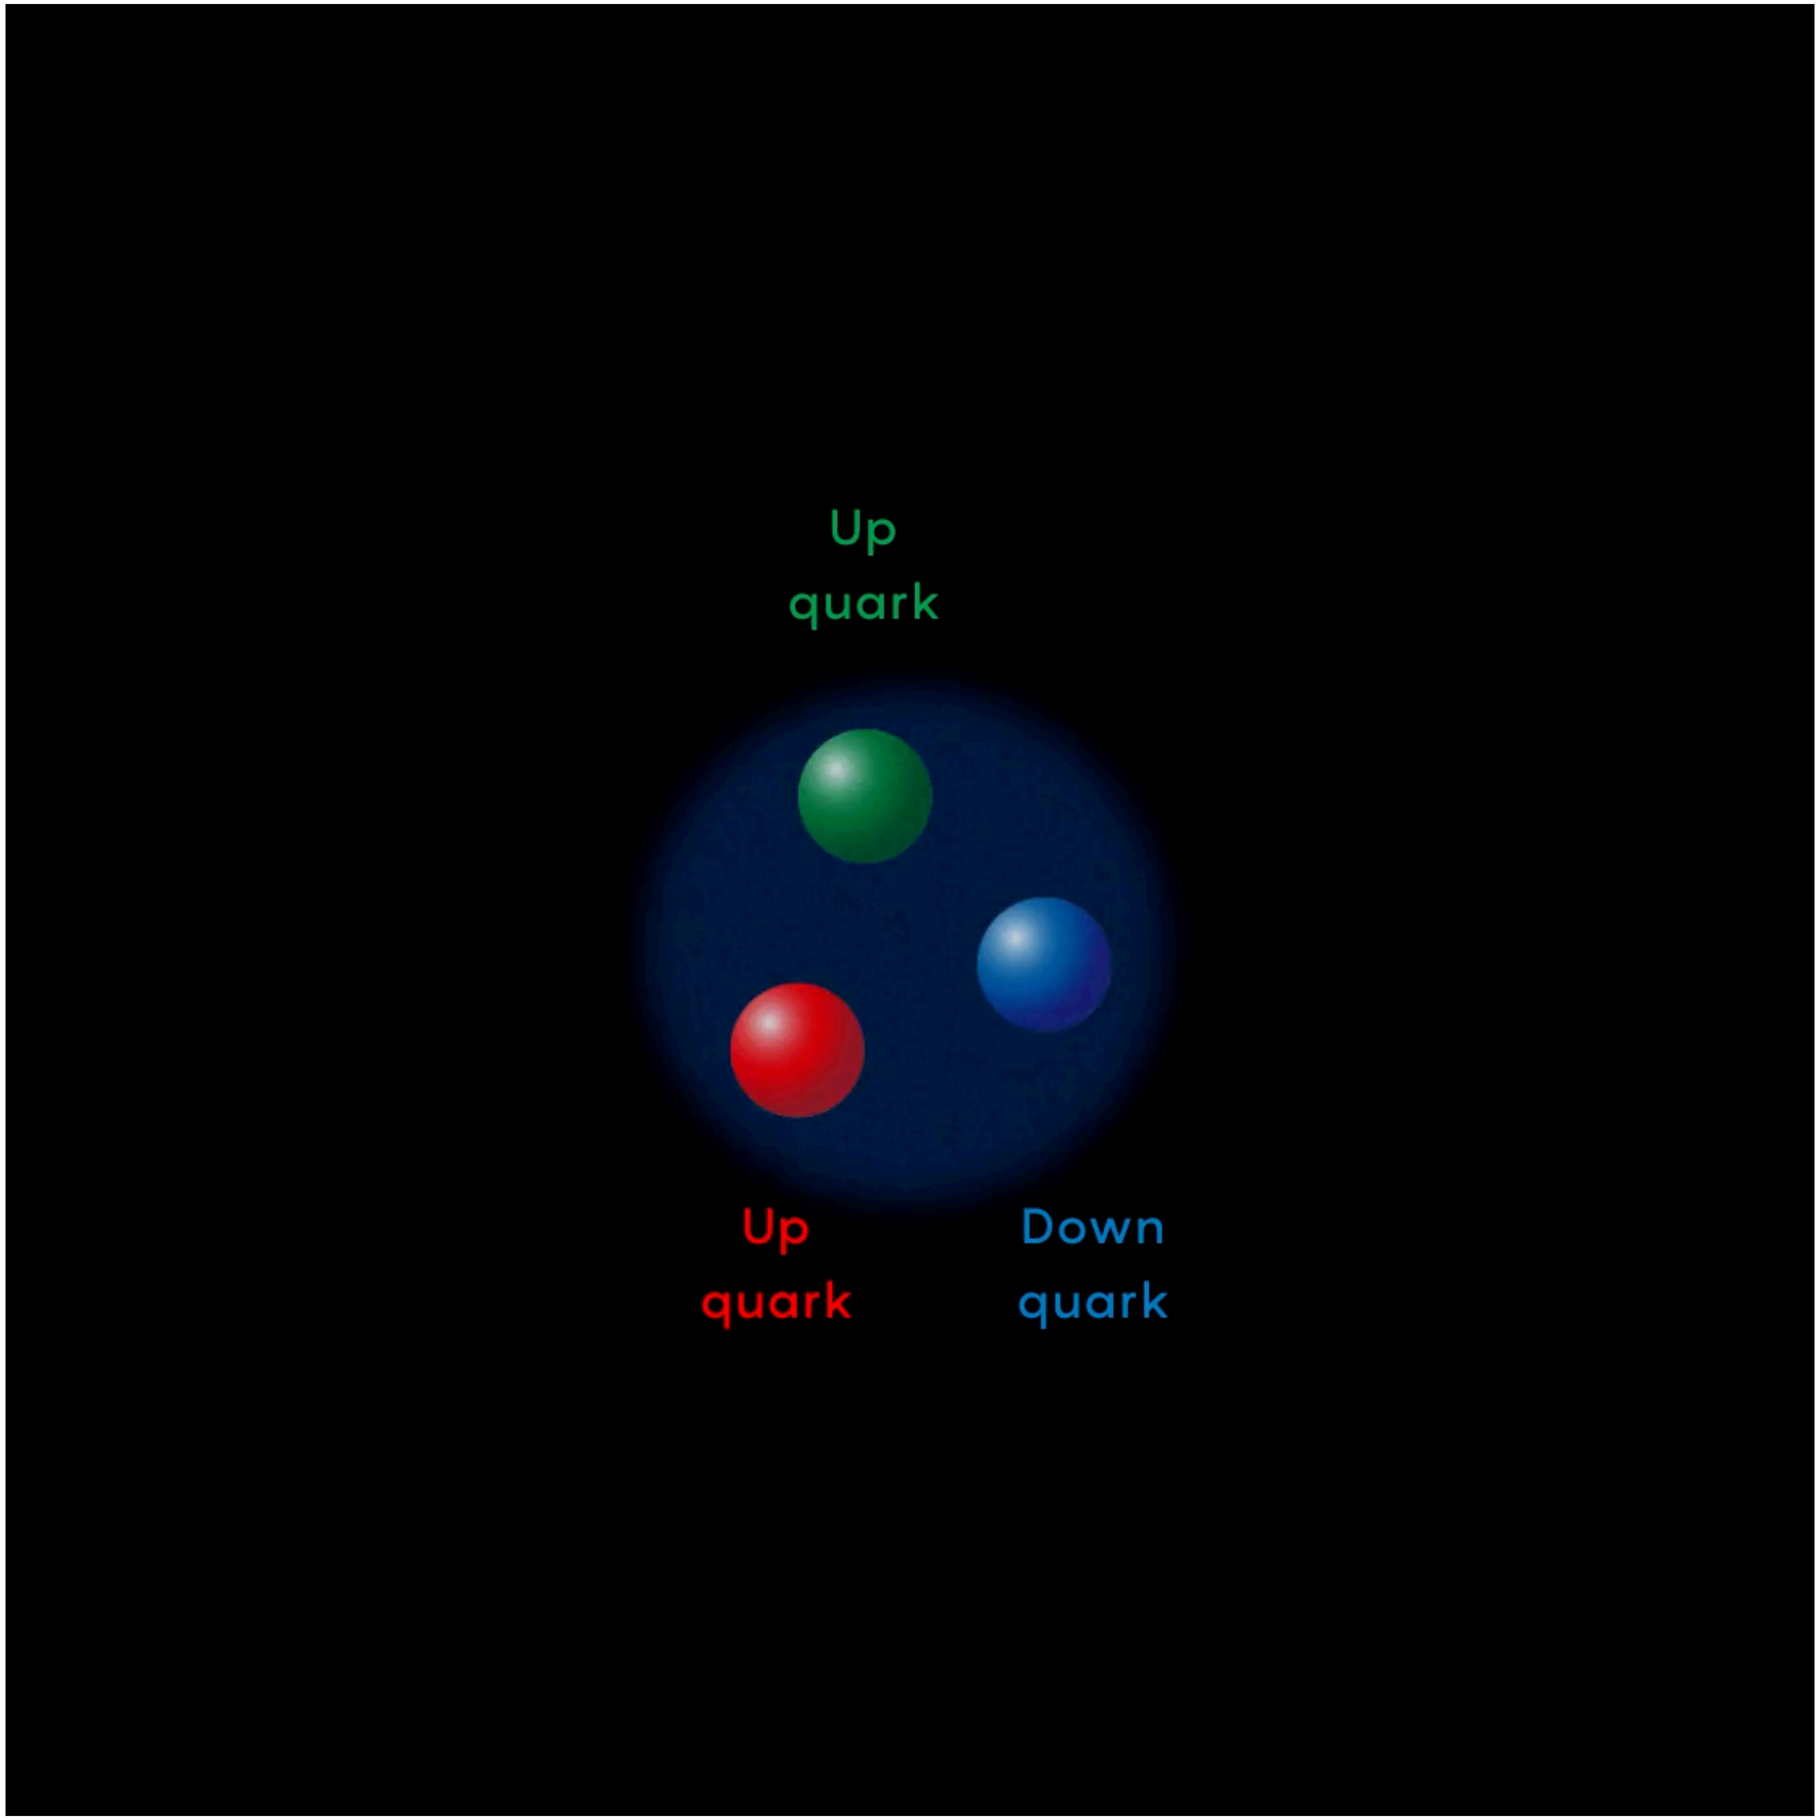

Charm in the proton

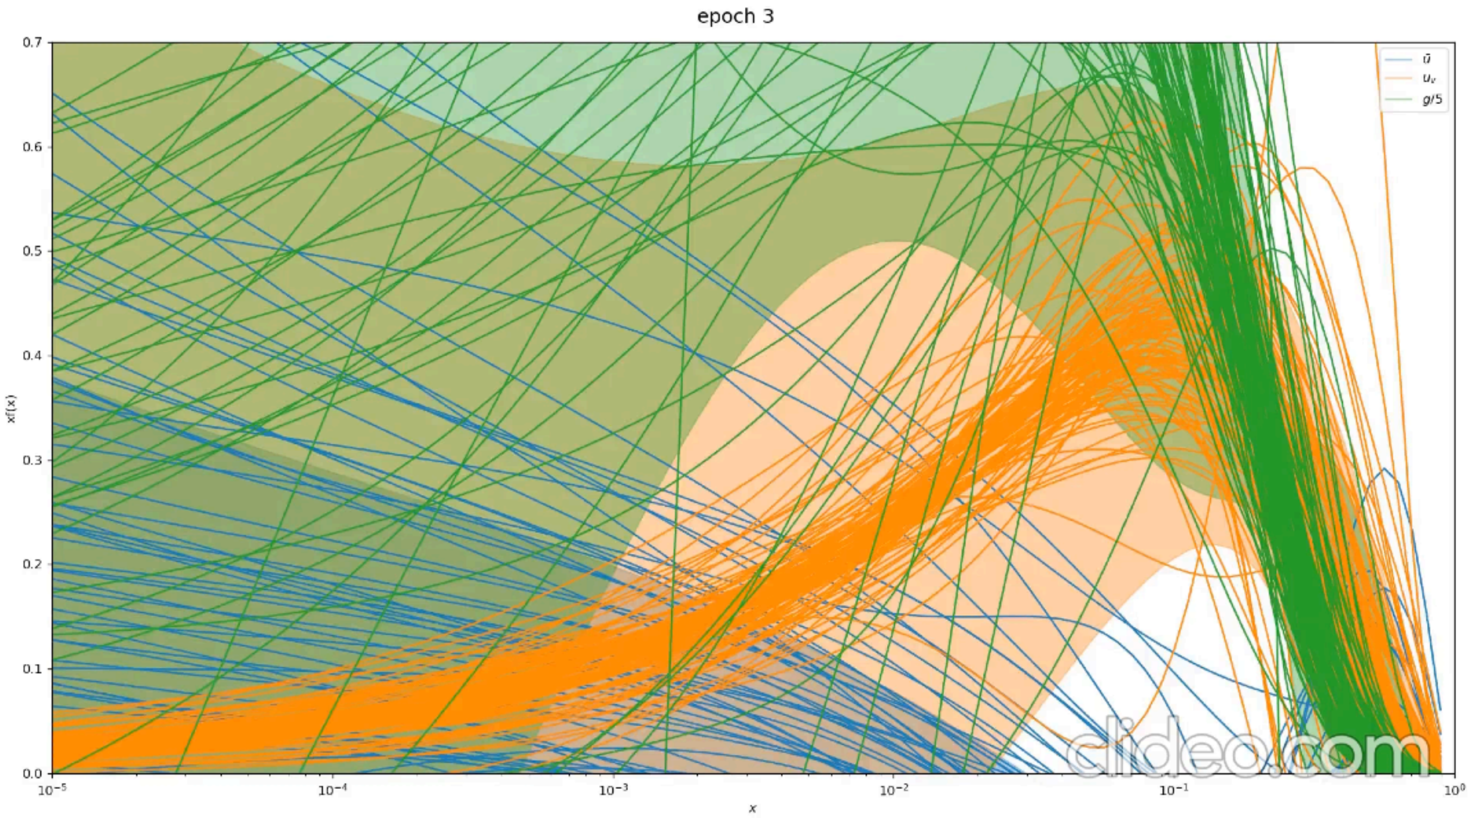

x: fraction of energy carried by proton constituents

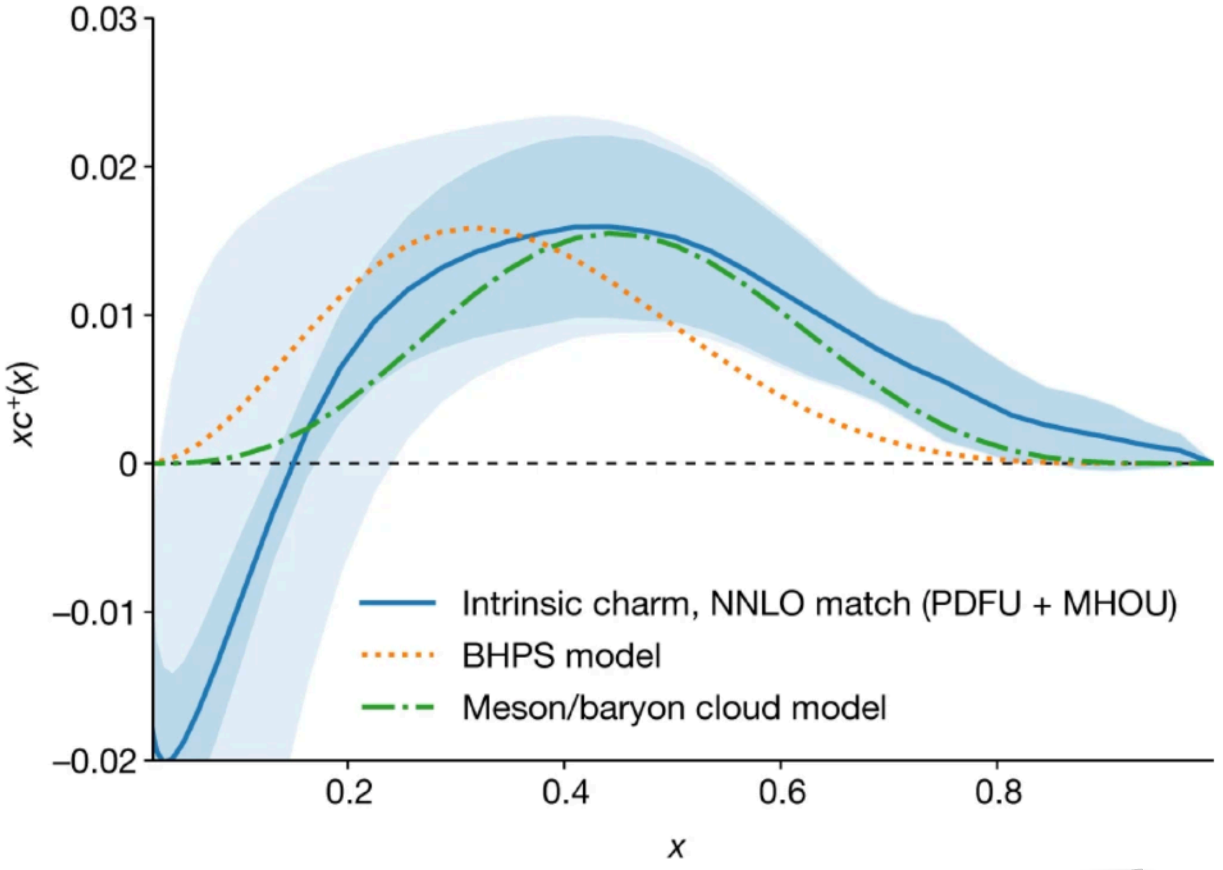

- **Machine learning** for unbiased regression of proton structure from data

- **Global analysis** of collider data with state-of-the-art QCD theory: evidence for charm in the proton

Article | Open access | Published: 17 August 2022

Evidence for intrinsic charm quarks in the proton

The NNPDF Collaboration

Nature 608, 483–487 (2022) | Cite this article

67k Accesses | 84 Citations | 368 Altmetric | Metrics

Abstract

The theory of the strong force, quantum chromodynamics, describes the proton in terms of quarks and gluons. The proton is a state of two up quarks and one down quark bound by gluons, but quantum theory predicts that in addition there is an infinite number of quark–antiquark pairs. Both light and heavy quarks, whose mass is respectively smaller or bigger than the mass of the proton, are revealed inside the proton in high-energy collisions. However, it is unclear whether heavy quarks also exist as a part of the proton wavefunction, which is determined by non-perturbative dynamics and accordingly unknown: so-called

Associated content

Evidence at last that the proton has intrinsic charm

Ramona Vogt
Nature | News & Views | 17 Aug 2022

Do protons have intrinsic charm? New evidence suggests yes

Benjamin Thompson & Nick Petrić Howe
Nature | Nature Podcast | 17 Aug 2022

Ball, …, JR (Nature, 2022)

Charm in the proton

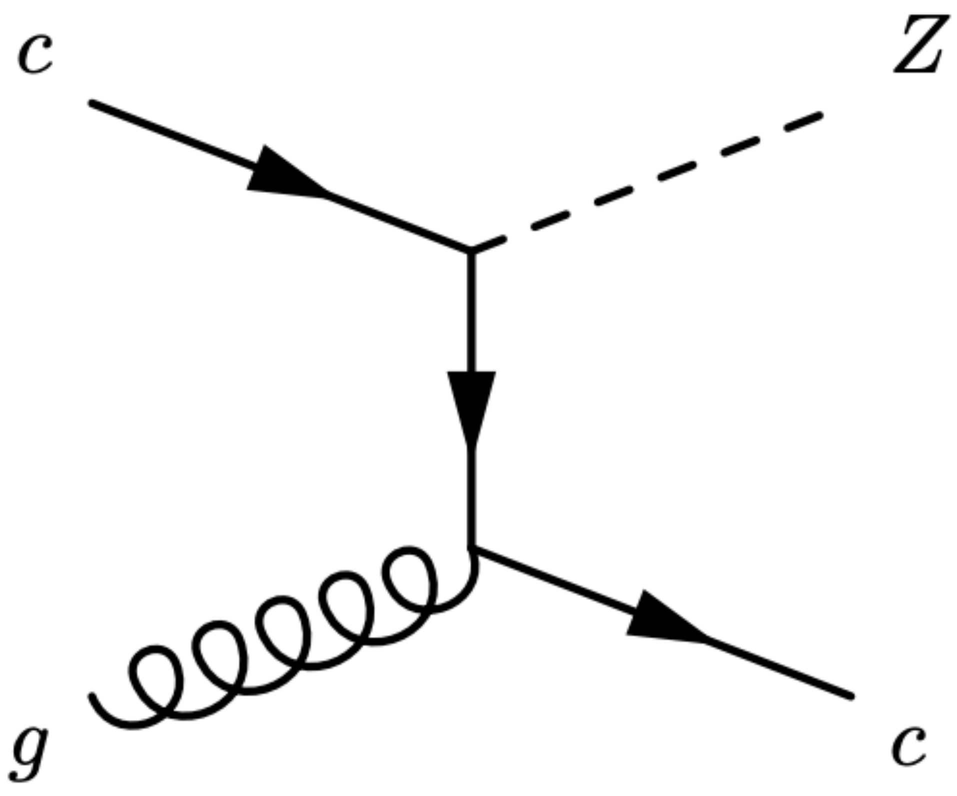

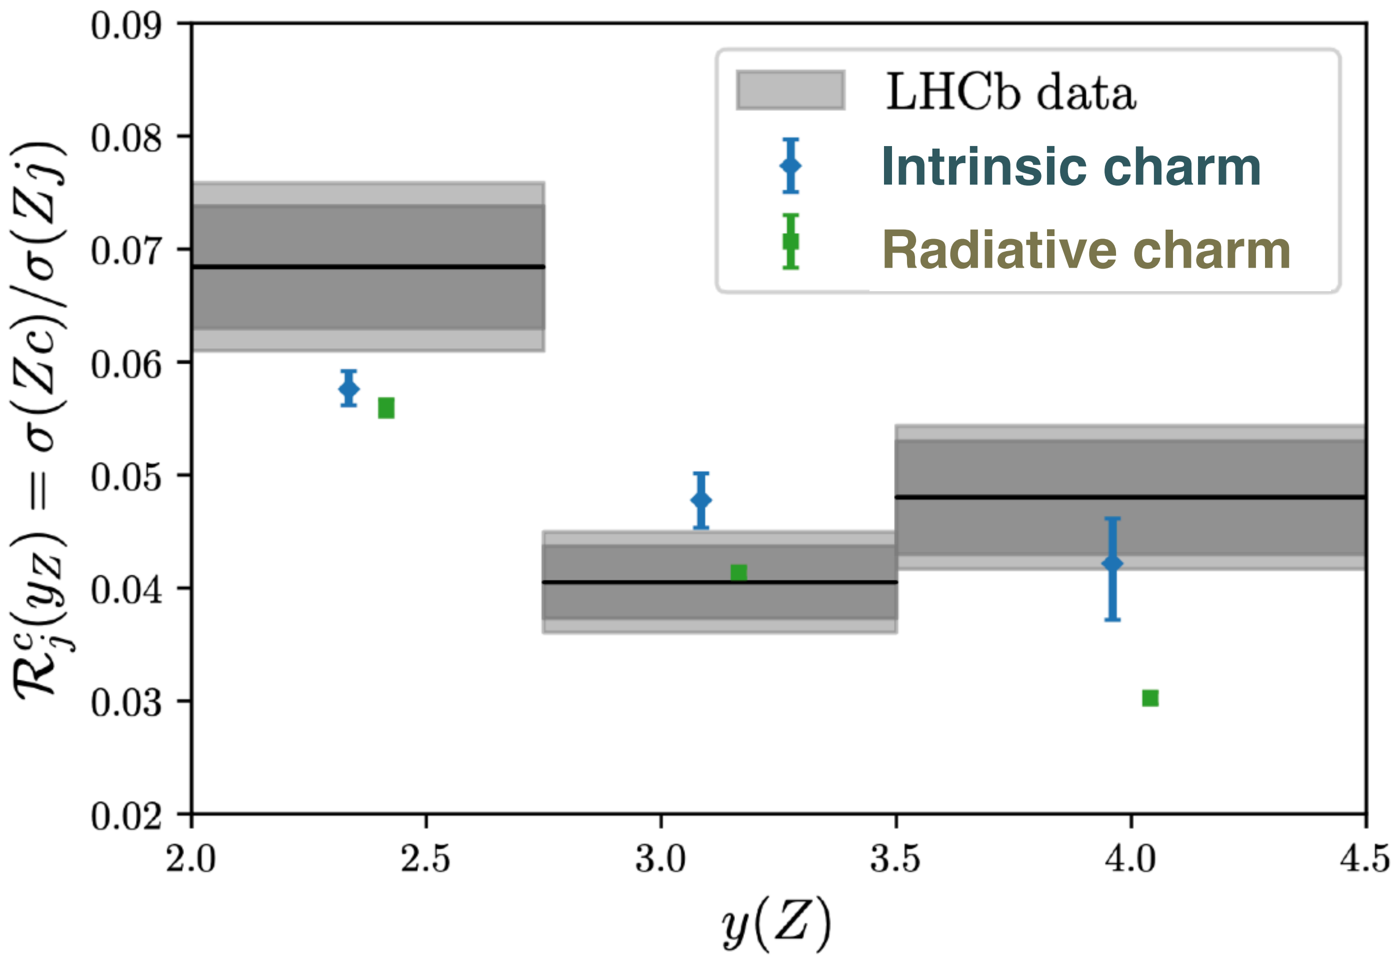

Measurements from the LHCb experiment at CERN indicate a **strong preference** for charm quarks in the proton

Key insights for our understanding of Quantum Chromodynamics

Implications for ultra-high energy astroparticle physics

AI for Optimally Sensitive Inference

The inference problem

- HEP experiments performs **searches** at the detector/reconstruction level
- **Measurements** are typically unfolded to particle- or parton-level binned distributions, *theory-friendly* quantities for their interpretation
- **Sensitivity loss** from unbinned detector data to binned parton-level cross-sections
- Solution: **train ML surrogates on simulations** to construct a **forward model** to carry our parameter inference with unbinned detector data

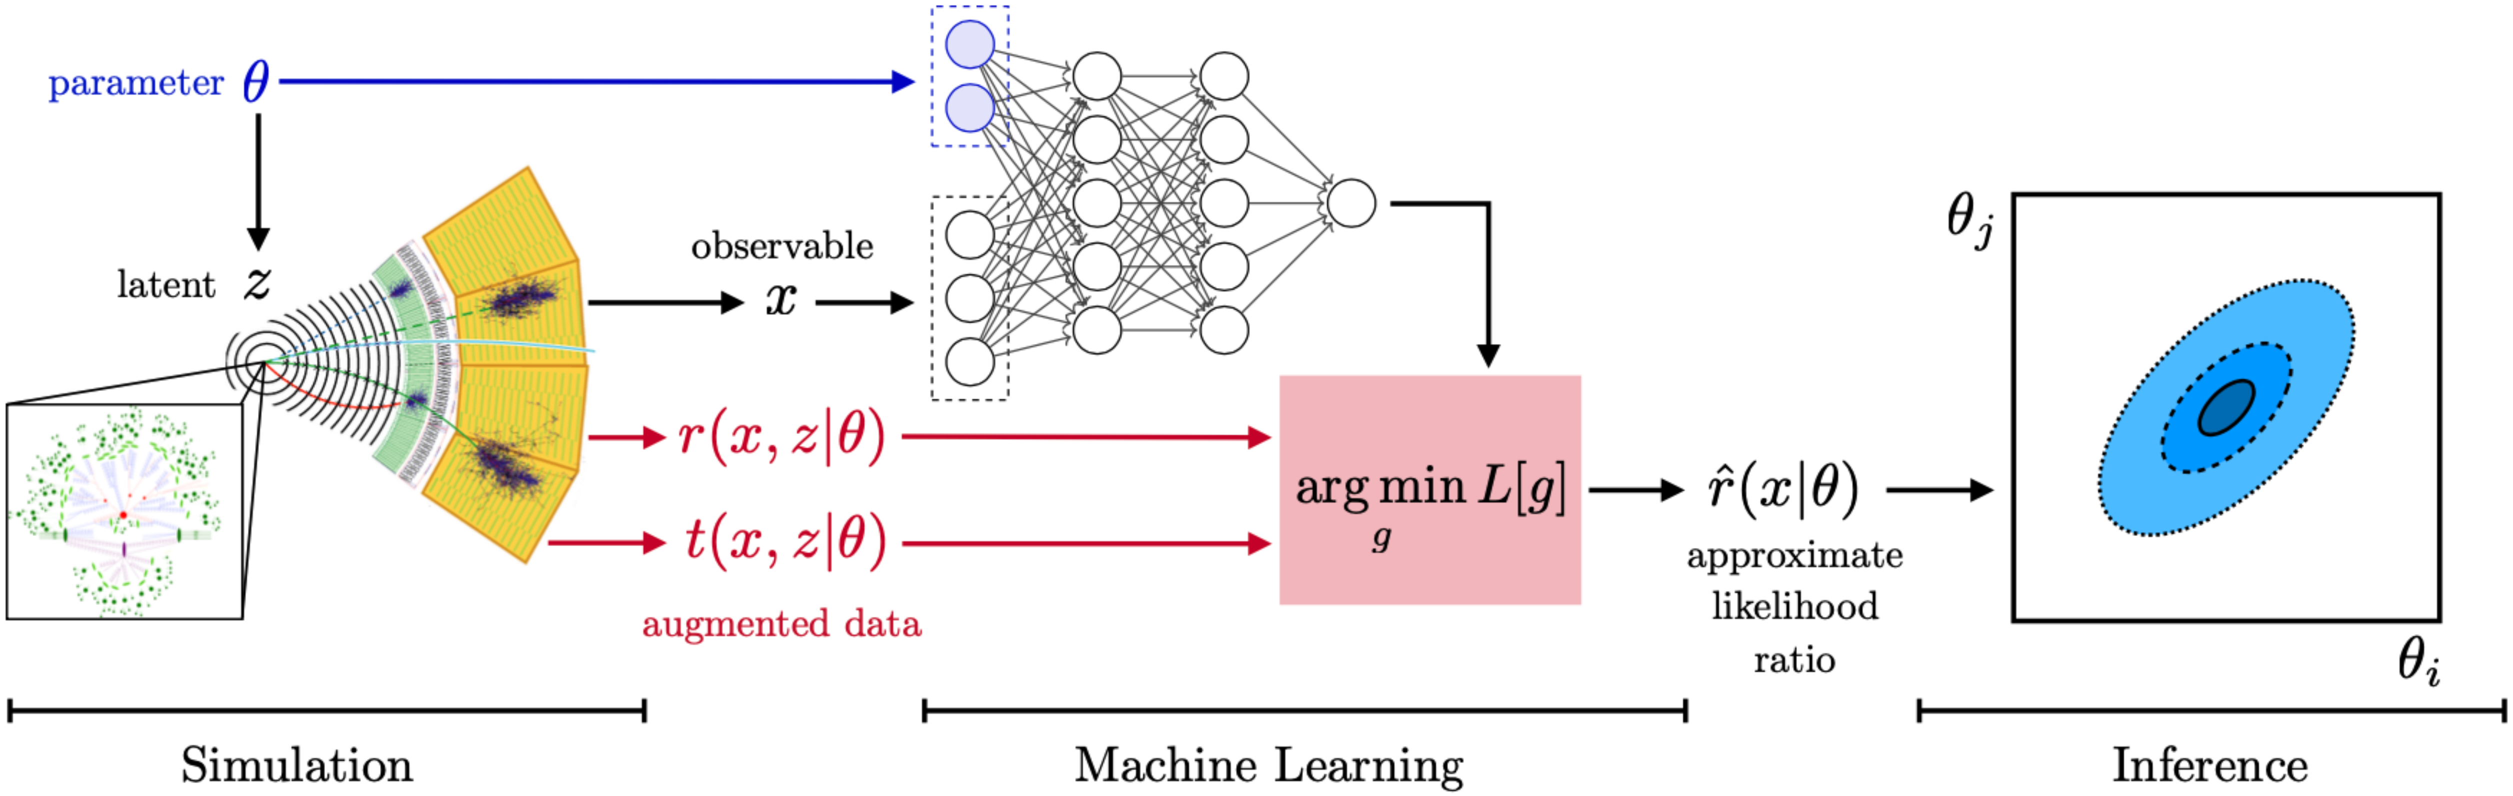

Neural Simulation-Based Inference

NSBI: learn from simulations mappings such as

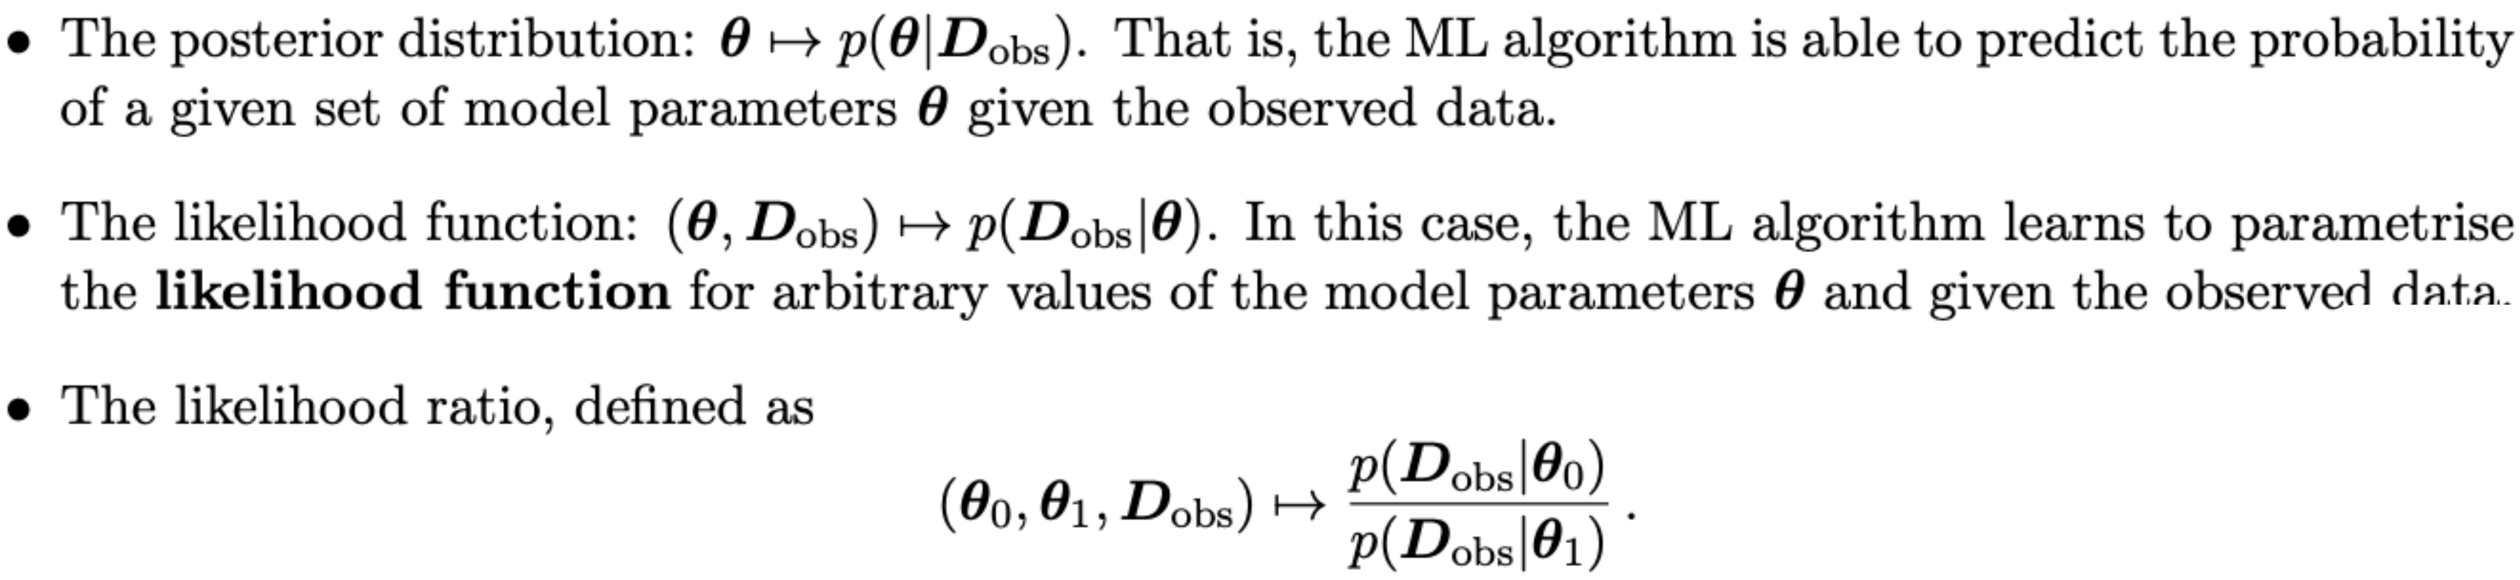

- The posterior distribution: $\boldsymbol{\theta} \mapsto p(\boldsymbol{\theta}|\boldsymbol{D}_{\text{obs}})$. That is, the ML algorithm is able to predict the probability of a given set of model parameters $\boldsymbol{\theta}$ given the observed data.
- The likelihood function: $(\boldsymbol{\theta}, \boldsymbol{D}_{\text{obs}}) \mapsto p(\boldsymbol{D}_{\text{obs}}|\boldsymbol{\theta})$. In this case, the ML algorithm learns to parametrise the **likelihood function** for arbitrary values of the model parameters $\boldsymbol{\theta}$ and given the observed data.
- The likelihood ratio, defined as

$$(\boldsymbol{\theta}_0, \boldsymbol{\theta}_1, \boldsymbol{D}_{\text{obs}}) \mapsto \frac{p(\boldsymbol{D}_{\text{obs}}|\boldsymbol{\theta}_0)}{p(\boldsymbol{D}_{\text{obs}}|\boldsymbol{\theta}_1)} .$$

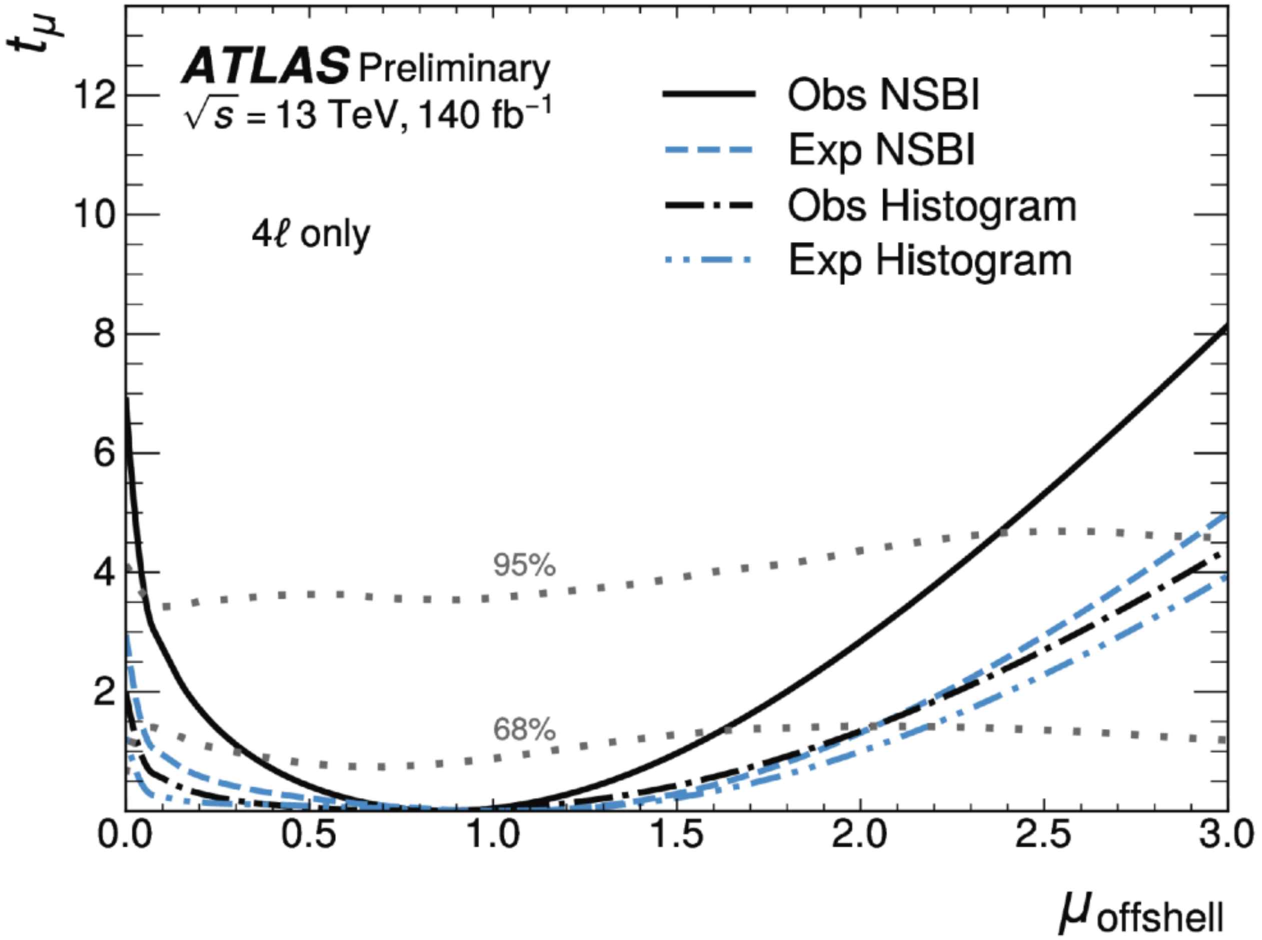

Higgs width from off-shell production with NSBI

Unbinned measurements approach **optimal sensitivity**: for a given dataset, you are guaranteed to use all available information

Neural Simulation-Based Inference

- NSBI involves learning complex functions in **very high-dimensional spaces**

All final-state features

$$L(\mathcal{D}|\boldsymbol{c}) = P(N_{\text{obs}}|\mathcal{L}\sigma(\boldsymbol{c})) \prod_{i=1}^{N_{\text{obs}}} \underbrace{\frac{1}{\sigma(\boldsymbol{c})}\frac{\mathrm{d}\sigma(\boldsymbol{x}|\boldsymbol{c})}{\mathrm{d}\boldsymbol{x}}\bigg|_{\boldsymbol{x}=\boldsymbol{x}_i}}_{p(\boldsymbol{x}|\boldsymbol{c})}$$

Dataset Likelihood

Parameters of Interest

Sum over events

Single Event Probability

detector-level simulations to train **ML surrogates**

$$R(\boldsymbol{x}|\boldsymbol{c}) = \frac{\mathrm{d}\sigma(\boldsymbol{x}|\boldsymbol{c})}{\mathrm{d}\sigma(\boldsymbol{x}|\boldsymbol{0})} \approx \hat{R}(\boldsymbol{x}|\boldsymbol{c}) = 1 + \sum_{A} c_A \Delta_A(\boldsymbol{x})$$

Crucial: linear expansion to decouple surrogate training

NSBI for Proton Structure

- Construct a **linear model** for the gluon PDF

$$f^{(g)}(x, Q_0, \boldsymbol{c}) = f^{(g)}_{\text{ref}}(x, Q_0) + \sum_{k=1}^{N} c_k \varphi_k^{(g)}(x)$$

- Learn **PDF variations** from event-by-event detector-level simulated MC sample

Lineal Model parameters

$$\omega(\boldsymbol{z}_i, \boldsymbol{c}) = \frac{f_{a_i}(x_{1i}; \mu_F, \boldsymbol{c})\, f_{b_i}(x_{2i}; \mu_F, \boldsymbol{c})}{f^{(\text{gen})}_{a_i}(x_{1i}; \mu_F)\, f^{(\text{gen})}_{b_i}(x_{2i}; \mu_F)}$$

$$\mathcal{L}_{\text{MSE}} = \sum_{\boldsymbol{c} \in \mathcal{V}} \sum_i w_{i,0} \left(\hat{R}(\boldsymbol{x}_i, \boldsymbol{c}) - \omega(\boldsymbol{z}_i, \boldsymbol{c}) \right)^2 \qquad \hat{R}(\boldsymbol{x}, \boldsymbol{c}) = 1 + \sum_A c_A \Delta_A(\boldsymbol{x})$$

ML surrogates

Include systematics errors

$$L(\mathcal{D}|\boldsymbol{c}, \boldsymbol{\nu}) = \text{Pois}(N_{\text{obs}} \,|\, \mathcal{L}(\boldsymbol{\nu})\, \sigma(\boldsymbol{c}, \boldsymbol{\nu})) \prod_{j=1} p(\boldsymbol{x}_j | \boldsymbol{c}, \boldsymbol{\nu}) \,,$$

- Use resulting forward for NSBI of proton structure

$$p(\boldsymbol{x}|\boldsymbol{c}, \boldsymbol{\nu}) = \frac{1}{\sigma(\boldsymbol{c}, \boldsymbol{\nu})} \frac{\mathrm{d}\sigma(\boldsymbol{x}|\boldsymbol{c}, \boldsymbol{\nu})}{\mathrm{d}\boldsymbol{x}} \,.$$

NSBI for Proton Structure

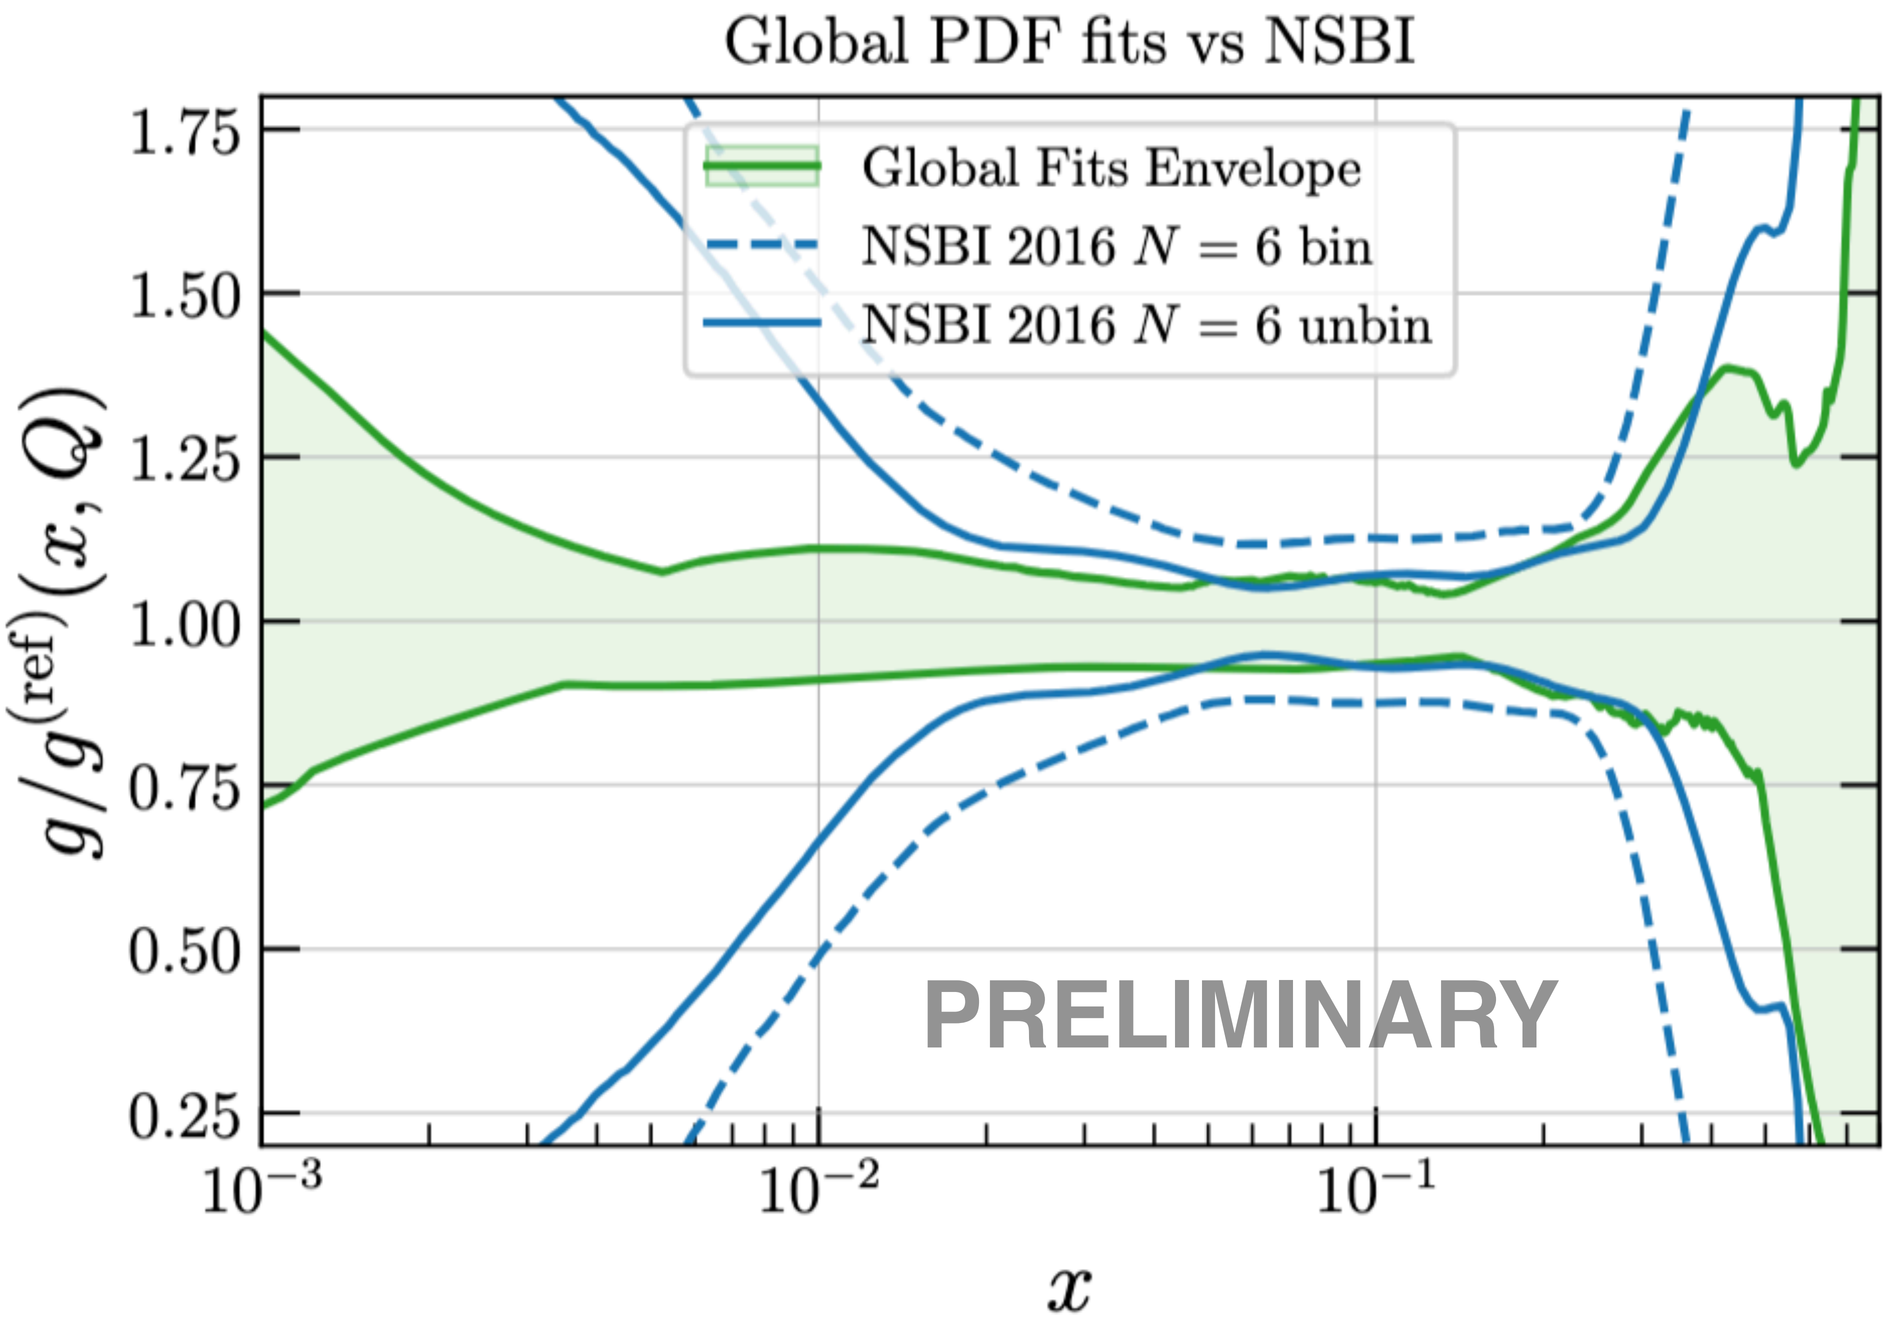

- Determination of the gluon PDF from top-quark pair production at LHC
- NSBI analysis: **comparable sensitivity to global fits** (very large dataset combinations)
- **Internal calibration of proton structure** within the LHC experiments

Barrue, Benato, Kaan Guven, Hammou, ter Hoeve, Krause, Li, Mantani, JR, Sanchez Cruz, Schofbeck, Ubiali, Wang (WIP)

ML and the Higgs Self-Interactions

Higgs Pair production

- Double Higgs production allows accessing **crucial components of the Higgs sector**:
  - ☑ Reconstructing the **electroweak symmetry breaking potential**
  - ☑ Probing the **Higgs self-interaction**
  - ☑ Assessing the **doublet nature** of the Higgs boson

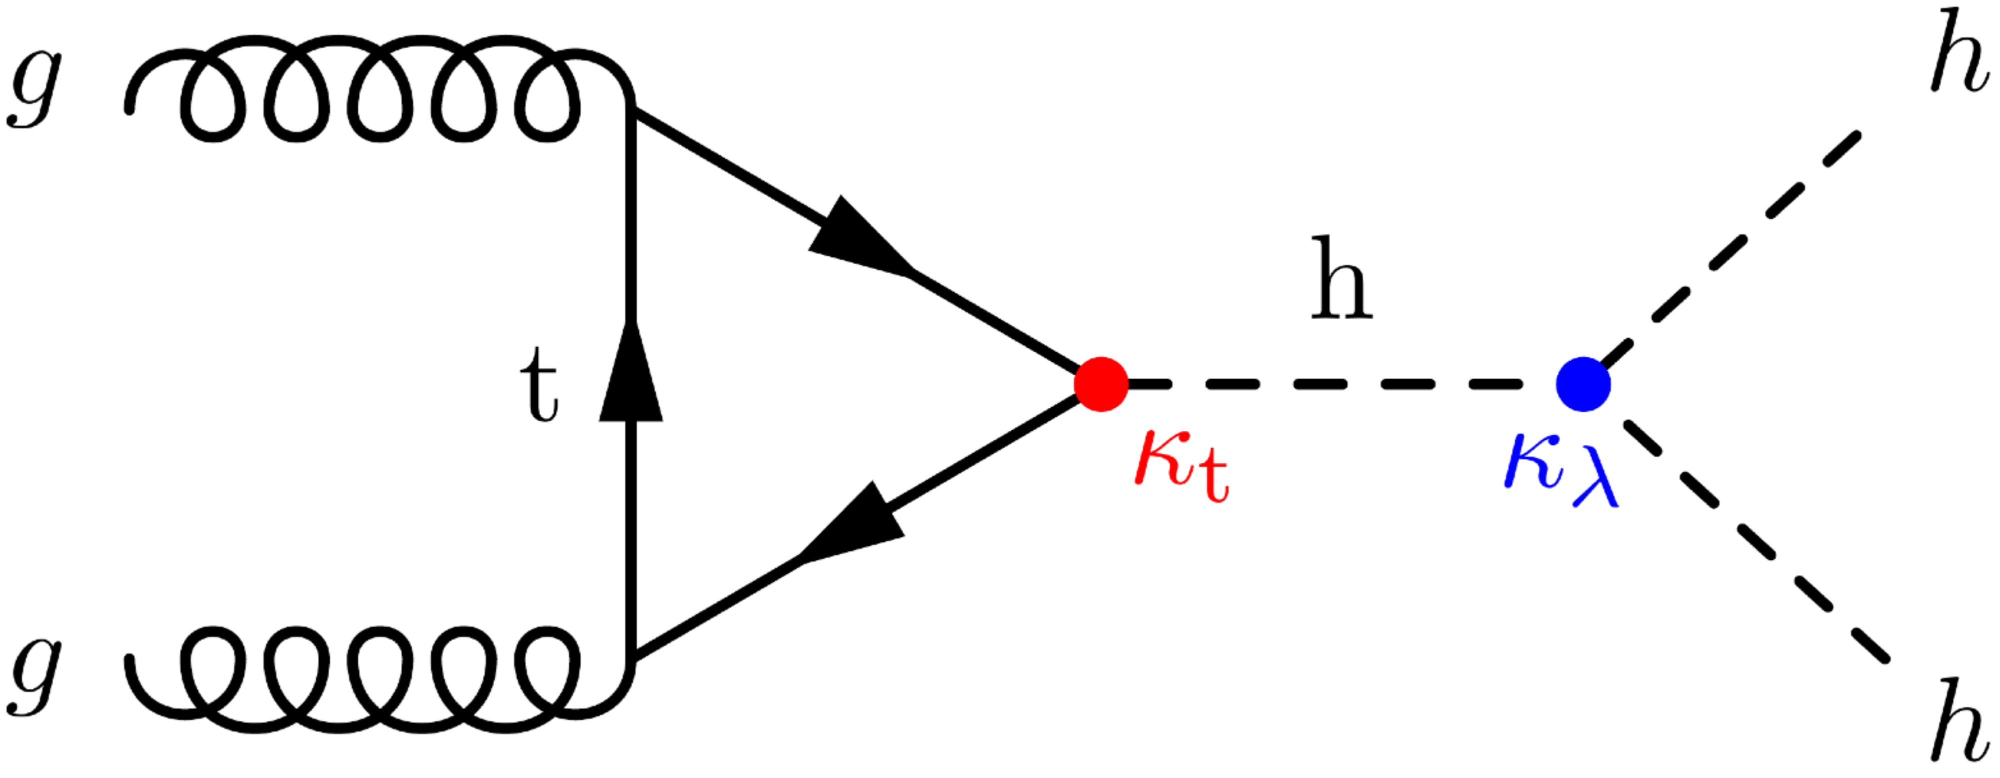

Higgs self-coupling: a unique new fundamental interaction (non-gauge & non-Yukawa!)

Higgs Pair production

- Double Higgs production allows accessing **crucial components of the Higgs sector**:
  - ☑ Reconstructing the **electroweak symmetry breaking potential**
  - ☑ Probing the **Higgs self-interaction**
  - ☑ Assessing the **doublet nature** of the Higgs boson

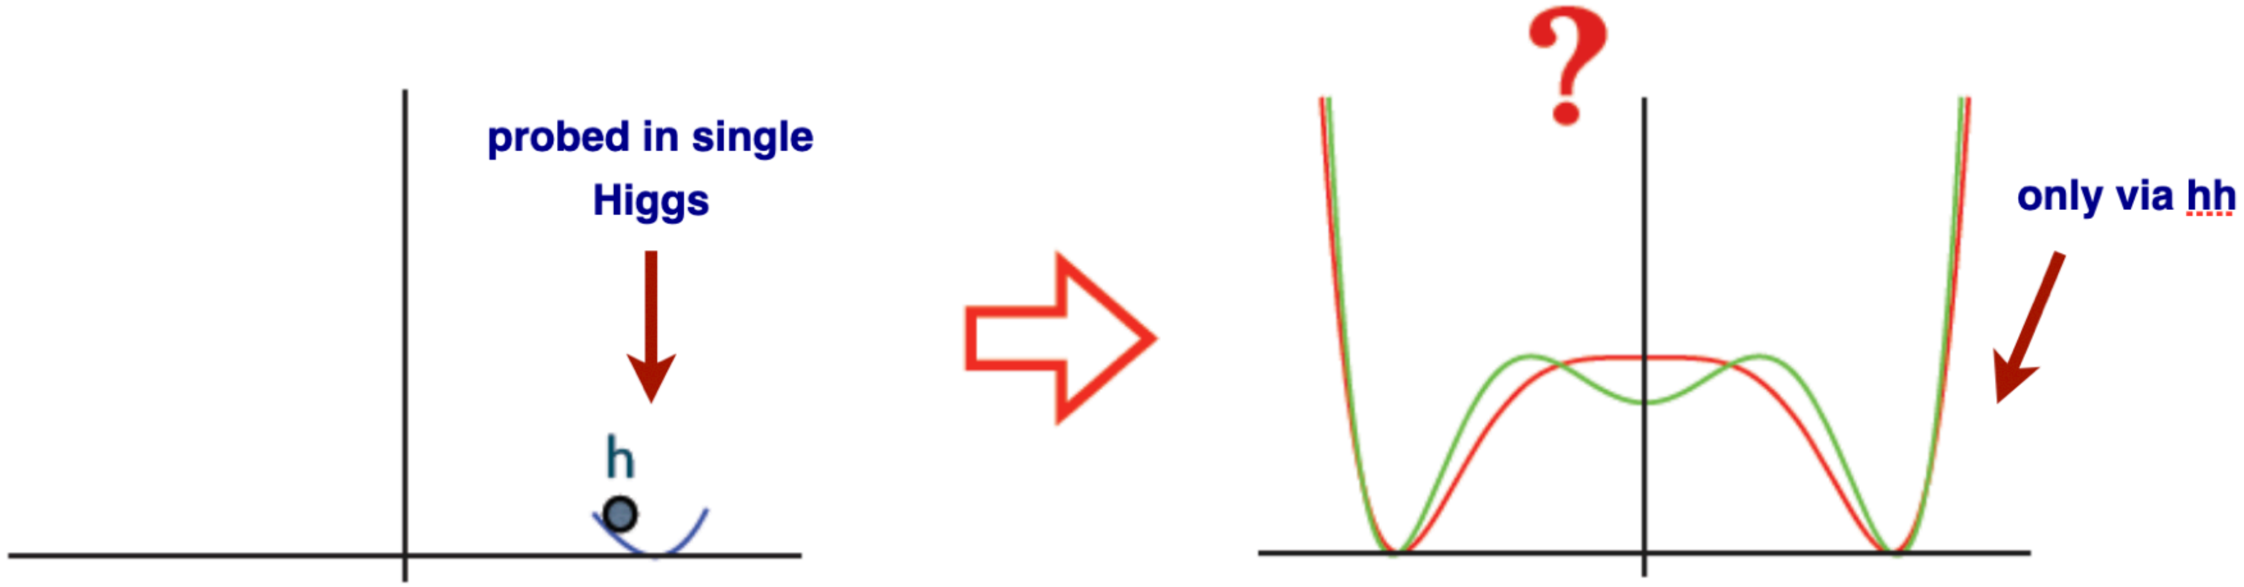

SM Higgs potential

$$V(h) = m_h^2 h^\dagger h + \frac{1}{2}\lambda(h^\dagger h)^2$$

Coleman-Weinberg potential

$$V(h) \to \frac{1}{2}\lambda(h^\dagger h)^2 \log\left[\frac{(h^\dagger h)}{m^2}\right]$$

Each possibility associated to **different EWSB mechanism**

Di-Higgs is almost here

- In the SM, **Higgs pair production rates are small**

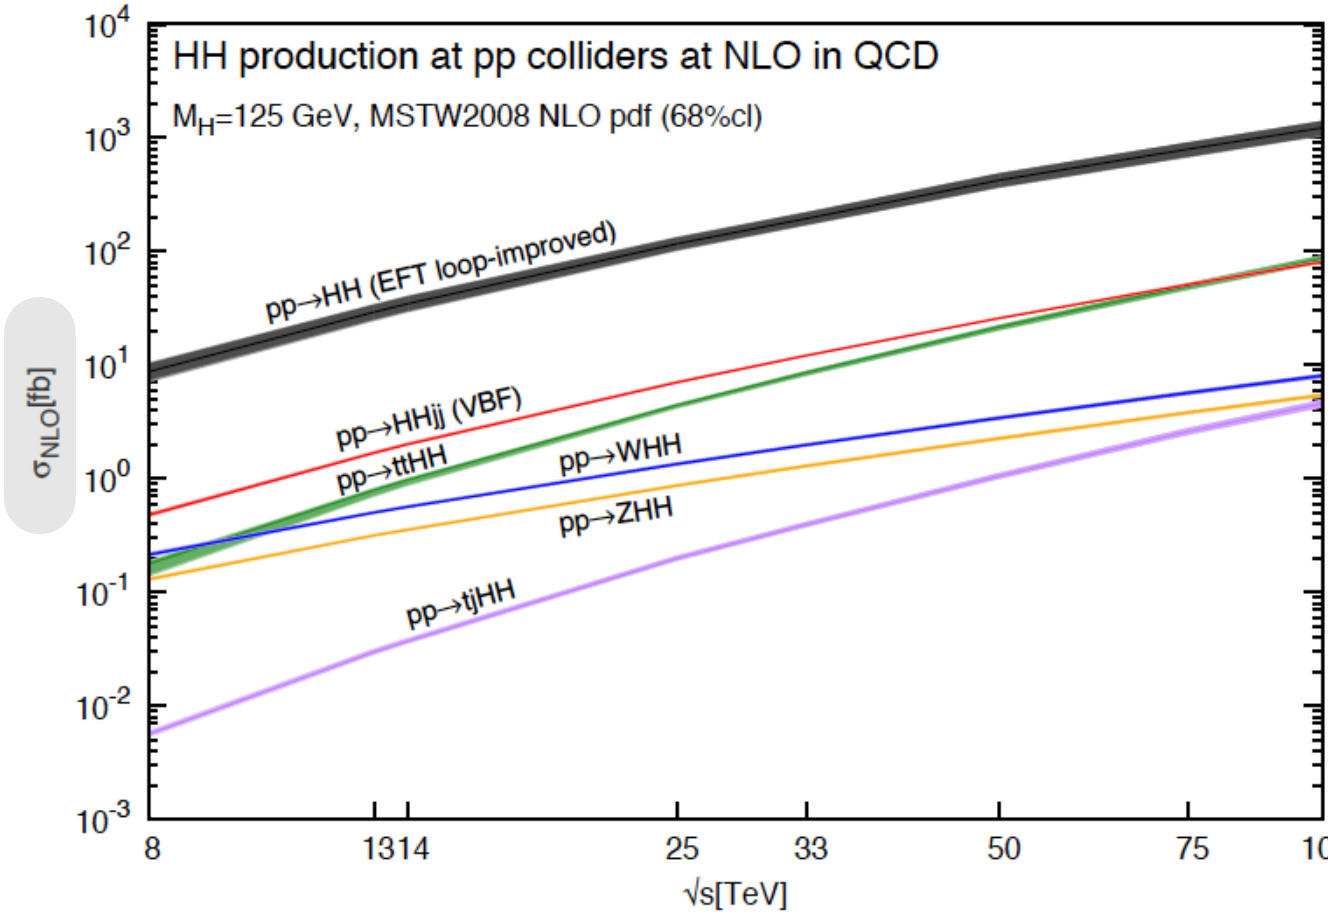

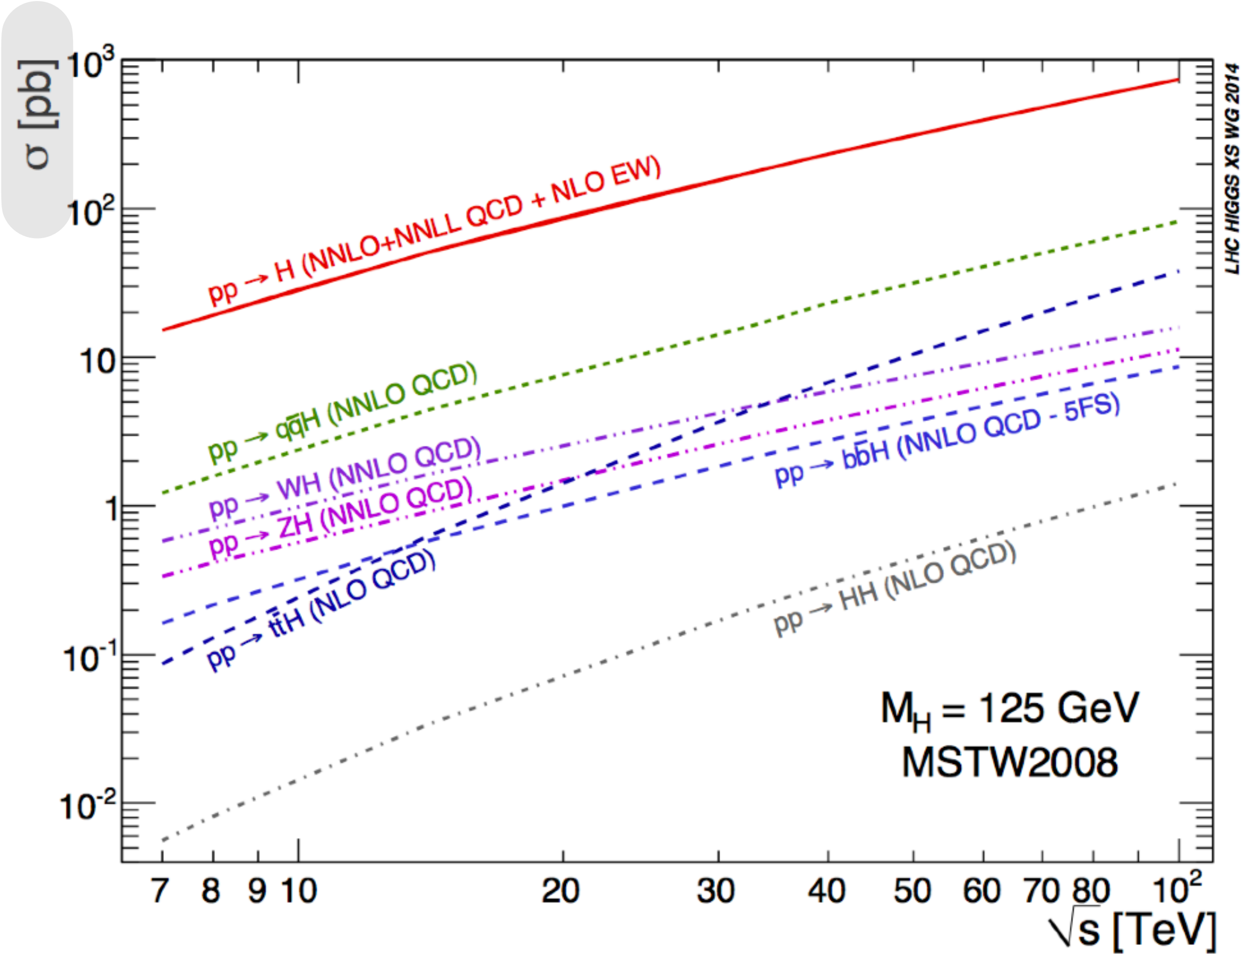

- **QCD backgrounds are large**, specially for **Higgs decays into a bottom-anti bottom pair**

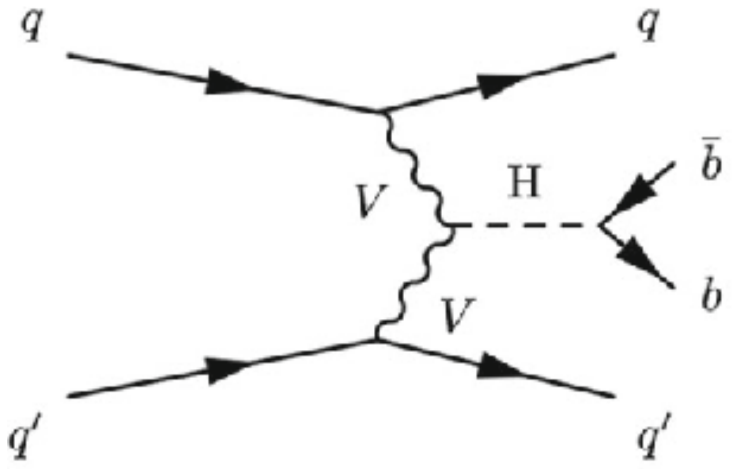

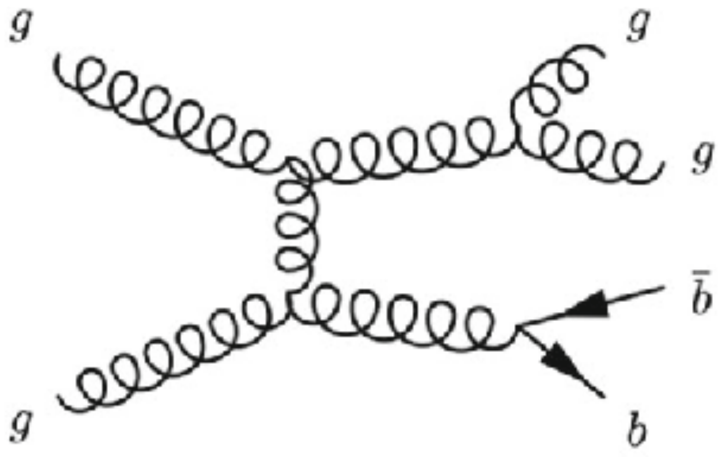

AI for classification

ML classifiers on simulations / control samples to enhance signal vs background discrimination

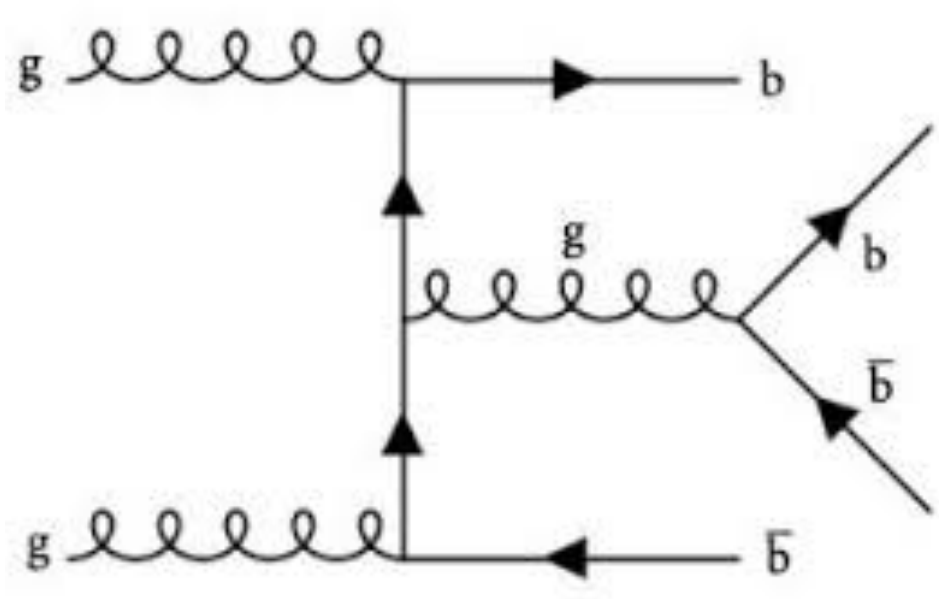

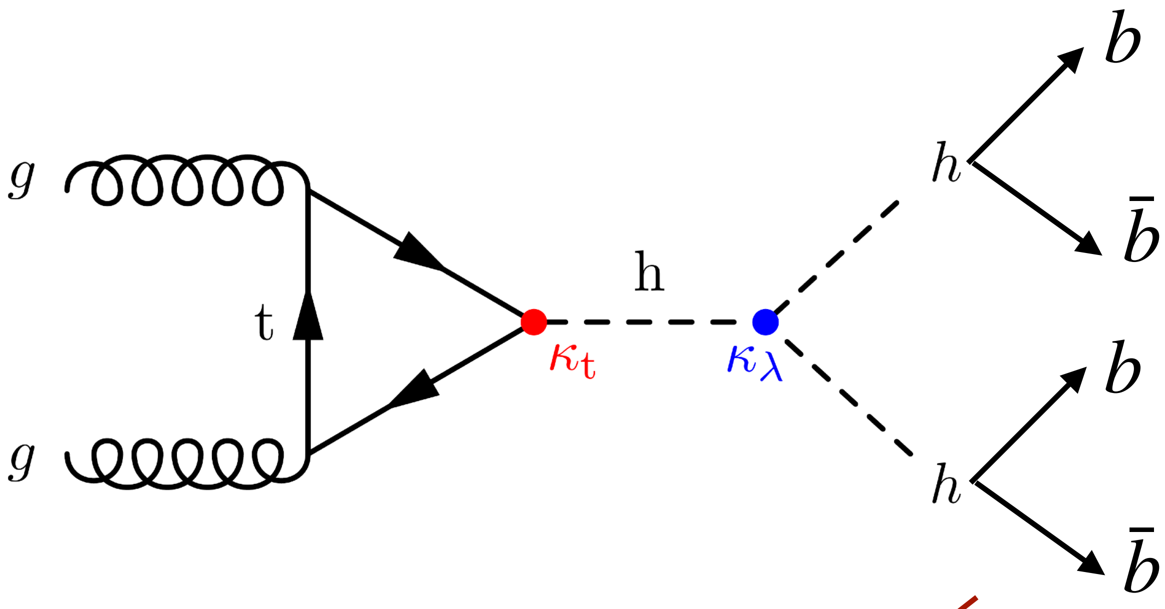

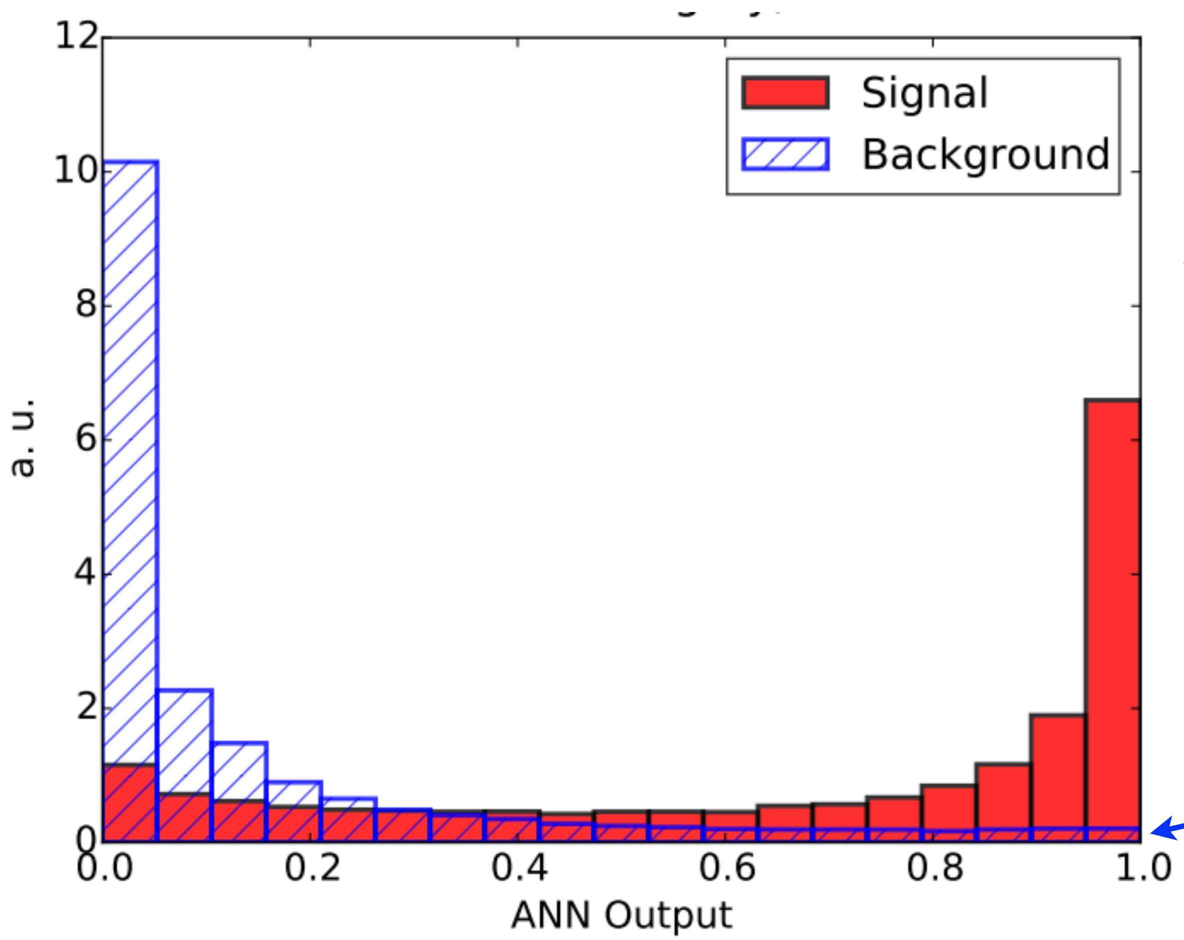

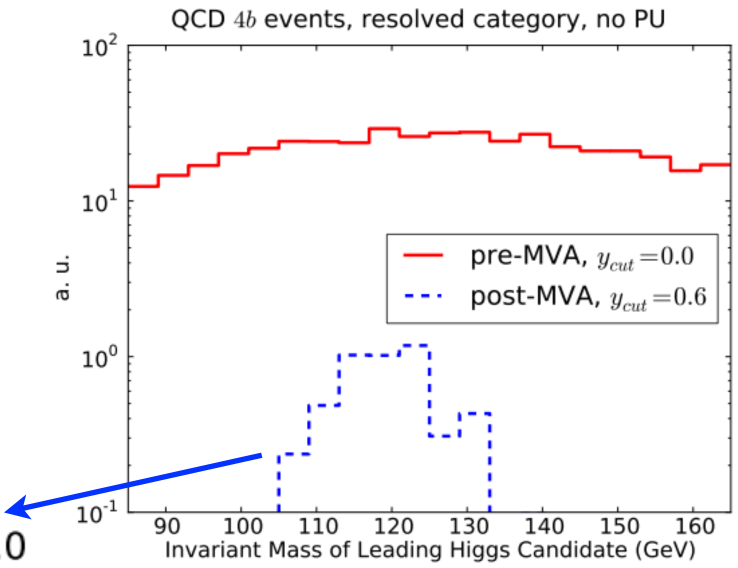

Di-Higgs is almost here

ML classification essential in the measurement of Higgs pair production

Eur. Phys. J. C (2016) 76:386
DOI 10.1140/epjc/s10052-016-4215-5

THE EUROPEAN
PHYSICAL JOURNAL C

CrossMark

Regular Article - Theoretical Physics

Boosting Higgs pair production in the $b\bar{b}b\bar{b}$ final state with multivariate techniques

J. Katharina Behr, Daniela Bortoletto, James A. Frost, Nathan P. Hartland, Cigdem Issever, Juan Rojo[a]
Physics Department, University of Oxford, 1 Keble Road, Oxford, UK

Received: 26 January 2016 / Accepted: 17 June 2016 / Published online: 8 July 2016

Abstract The measurement of Higgs pair production will be a cornerstone of the LHC program in the coming years. Double Higgs production provides a crucial window upon the mechanism of electroweak symmetry breaking and has a unique sensitivity to the Higgs trilinear coupling. We study the feasibility of a measurement of Higgs pair production in the $b\bar{b}b\bar{b}$ final state at the LHC. Our analysis is based on a combination of traditional cut-based methods with state-of-the-art multivariate techniques. We account for all relevant backgrounds, including the contributions from light and charm jet mis-identification, which are ultimately comparable in size to the irreducible $4b$ QCD background. We demonstrate the robustness of our analysis strategy in a high pileup environment. For an integrated luminosity of $\mathcal{L} = 3$ ab^{-1}, a signal significance of $S/\sqrt{B} \simeq 3$ is obtained, indicating that the $b\bar{b}b\bar{b}$ final state alone could allow for the observation of double Higgs production at the High Luminosity LHC.

tion will also be relevant to any future high-energy collider, either at a 100 TeV circular machine [11–14] or at a linear or circular electron–positron collider [15].

Analogously to single Higgs production [16], in the SM the dominant mechanism for the production of a pair of Higgs bosons at the LHC is gluon fusion (see [3,17] and references therein). For a center-of-mass energy of $\sqrt{s} = 14$ TeV, the next-to-next-to-leading order (NNLO) total cross section is approximately 40 fb [18], which is increased by a further few percent once next-to-next-to-leading logarithmic (NNLL) corrections are accounted for [19]. Feasibility studies in the case of a SM-like Higgs boson in the gluon-fusion channel at the LHC have been performed for different final states, including $b\bar{b}\gamma\gamma$ [20–22], $b\bar{b}\tau^+\tau^-$ [23–26], $b\bar{b}W^+W^-$ [25,27] and $b\bar{b}b\bar{b}$ [21,23,25,28,29]. While these studies differ in their quantitative conclusions, the consistent picture emerges that the ultimate precision in the determination of the Higgs trilinear coupling λ requires the full inte-

- ML-based classifier combining **boosted and resolved topologies**

- We found that full HL-LHC luminosity required to **observe** $hh \rightarrow 4b$

$$\left(\frac{S}{\sqrt{B}}\right)_{\rm tot} \simeq 3.1\ (1.0), \quad \mathcal{L} = 3000\ (300)\ {\rm fb}^{-1},$$

extracting Higgs trilinear coupling remained very challenging, even at the HL-LHC, in this final state

- Recent progress in ML classifiers improve even the most optimistic projections

AI for classification

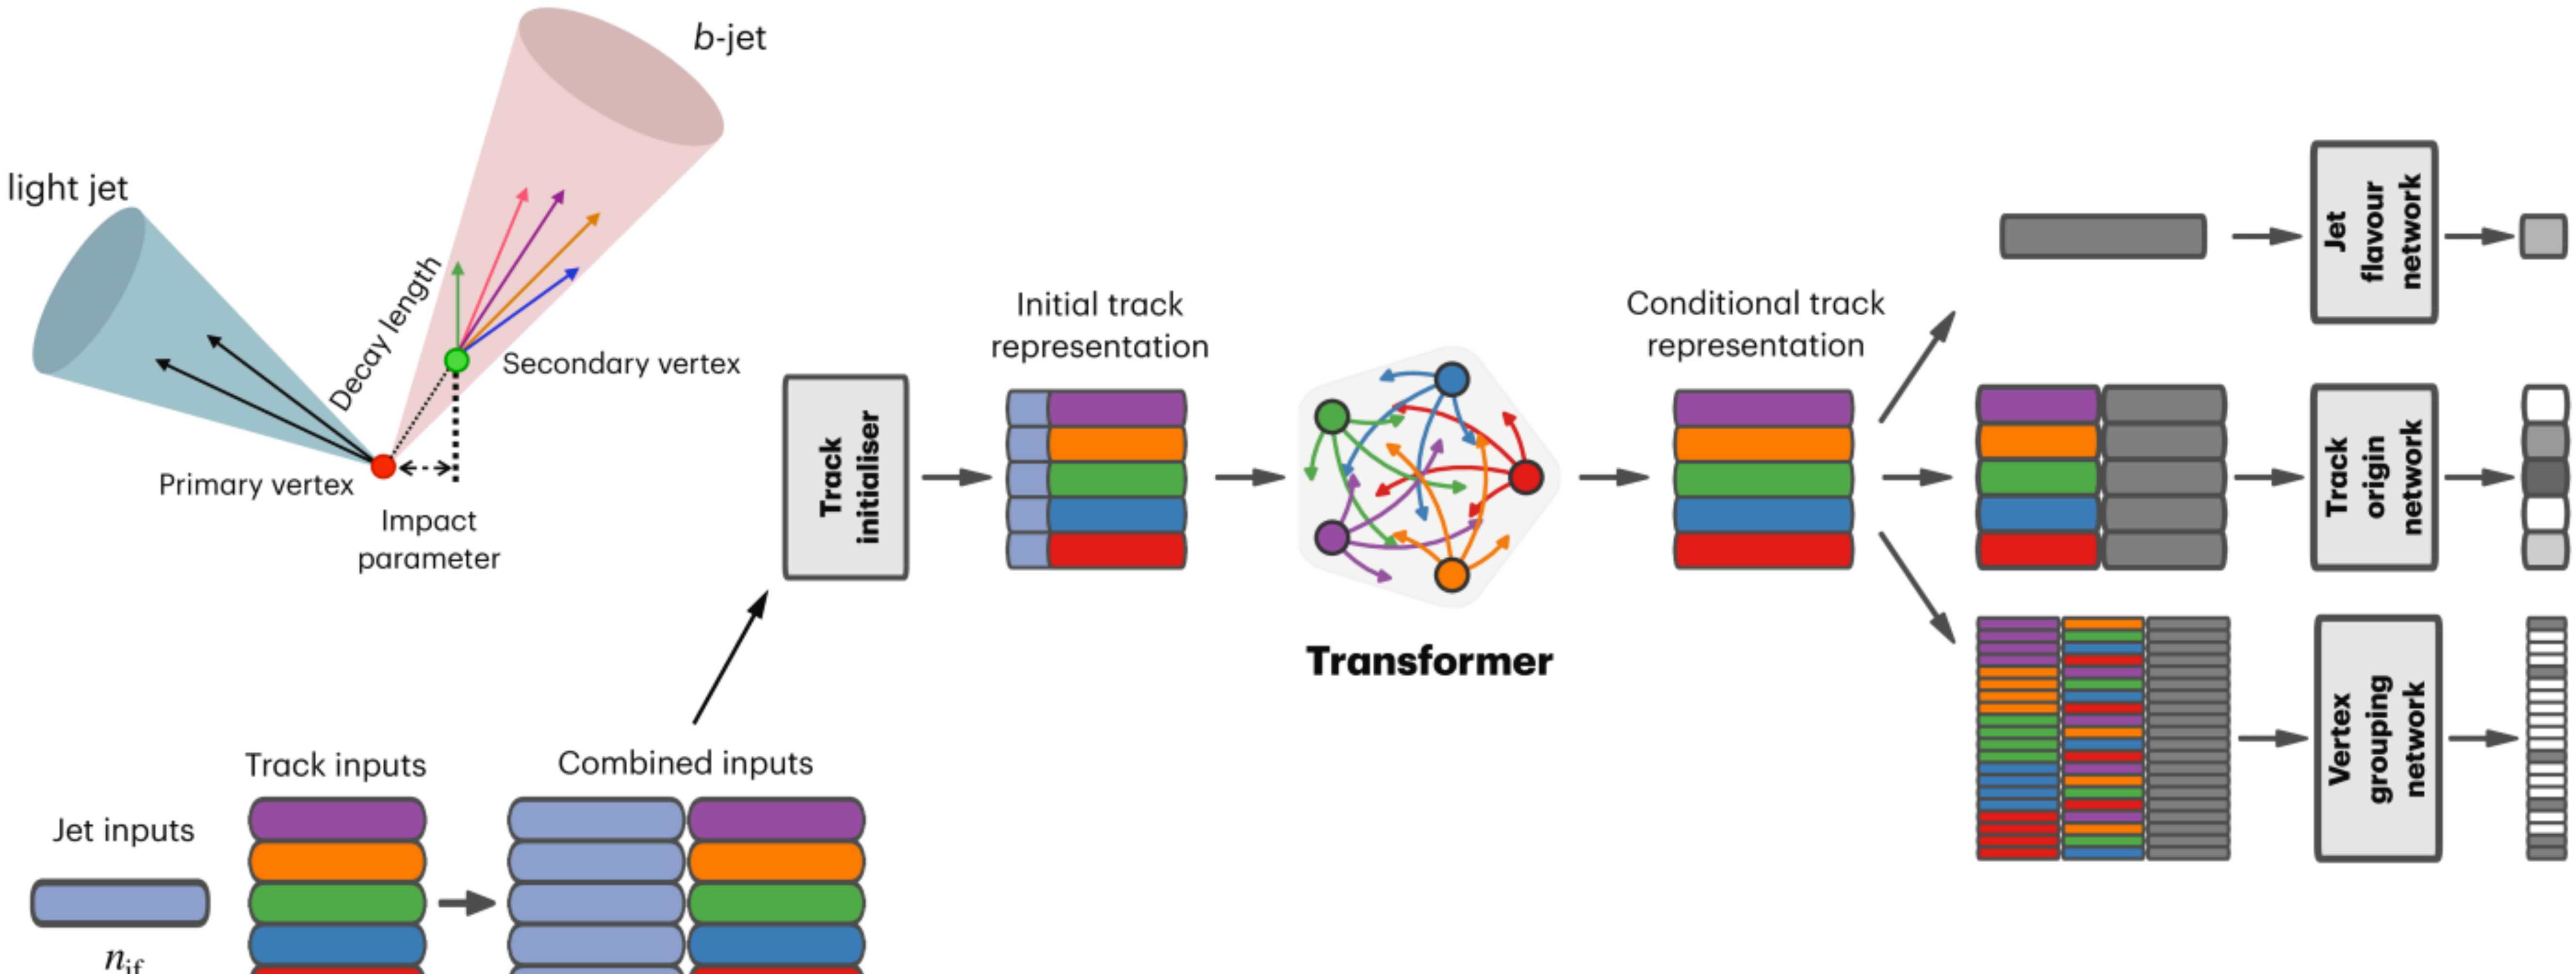

- **Graph Neural Networks** & **Transformers** enable greatly superior b-jet tagging at the LHC

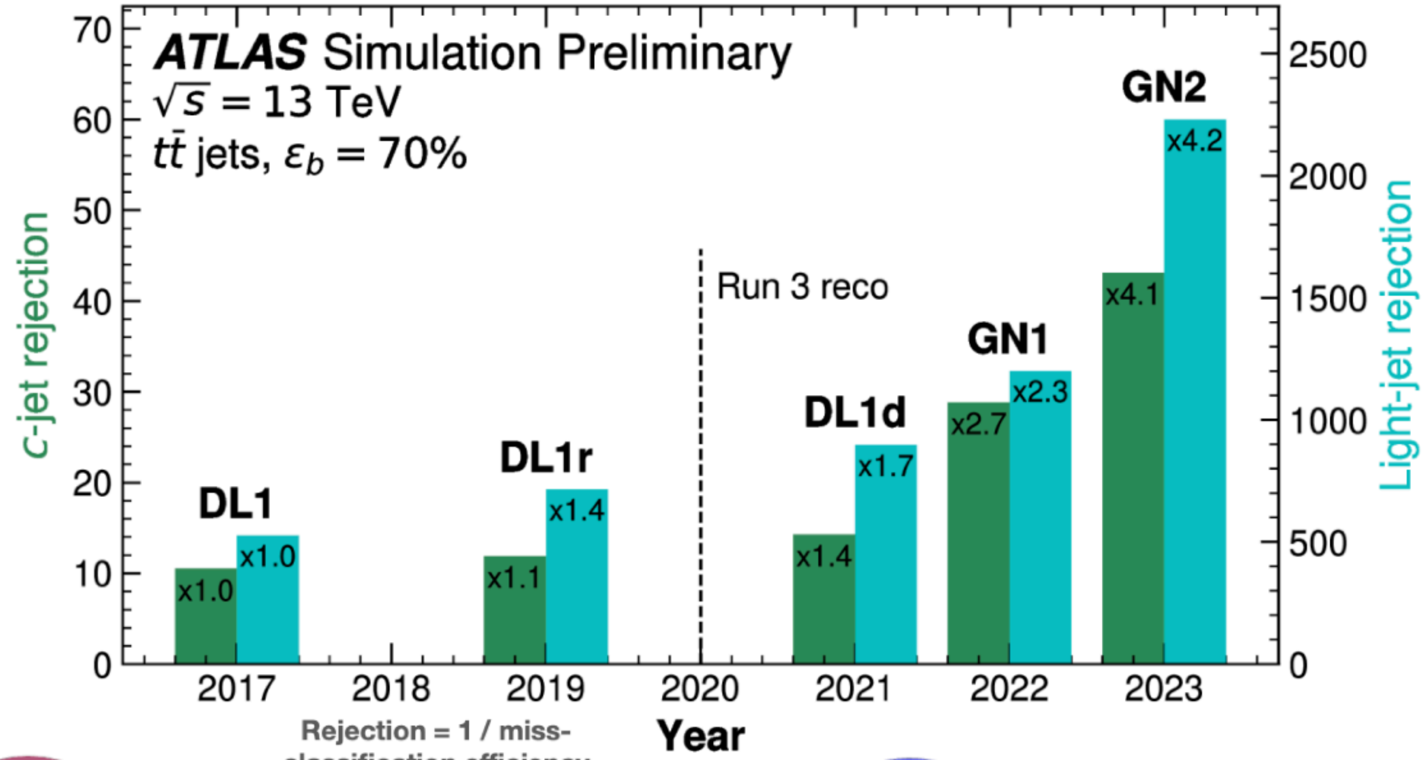

Rejection = 1 / miss-

ESPPU19 vs ESPPU26

📌 Do we need a **post-LHC collider** to measure the Higgs self-coupling?

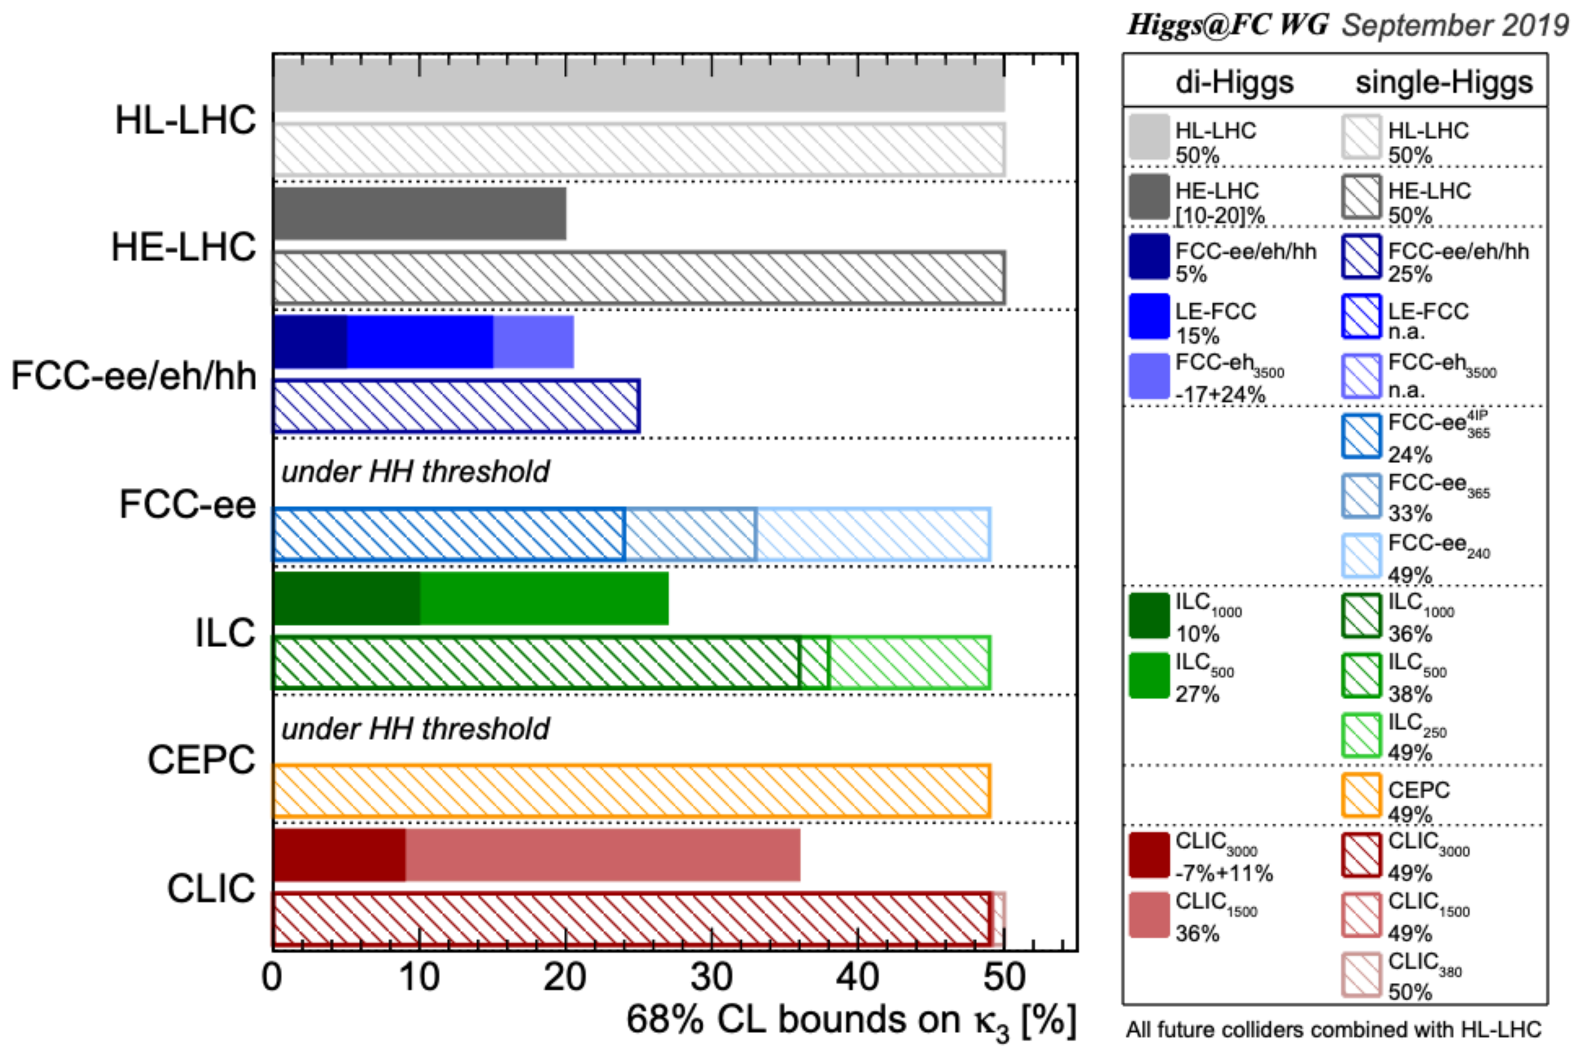

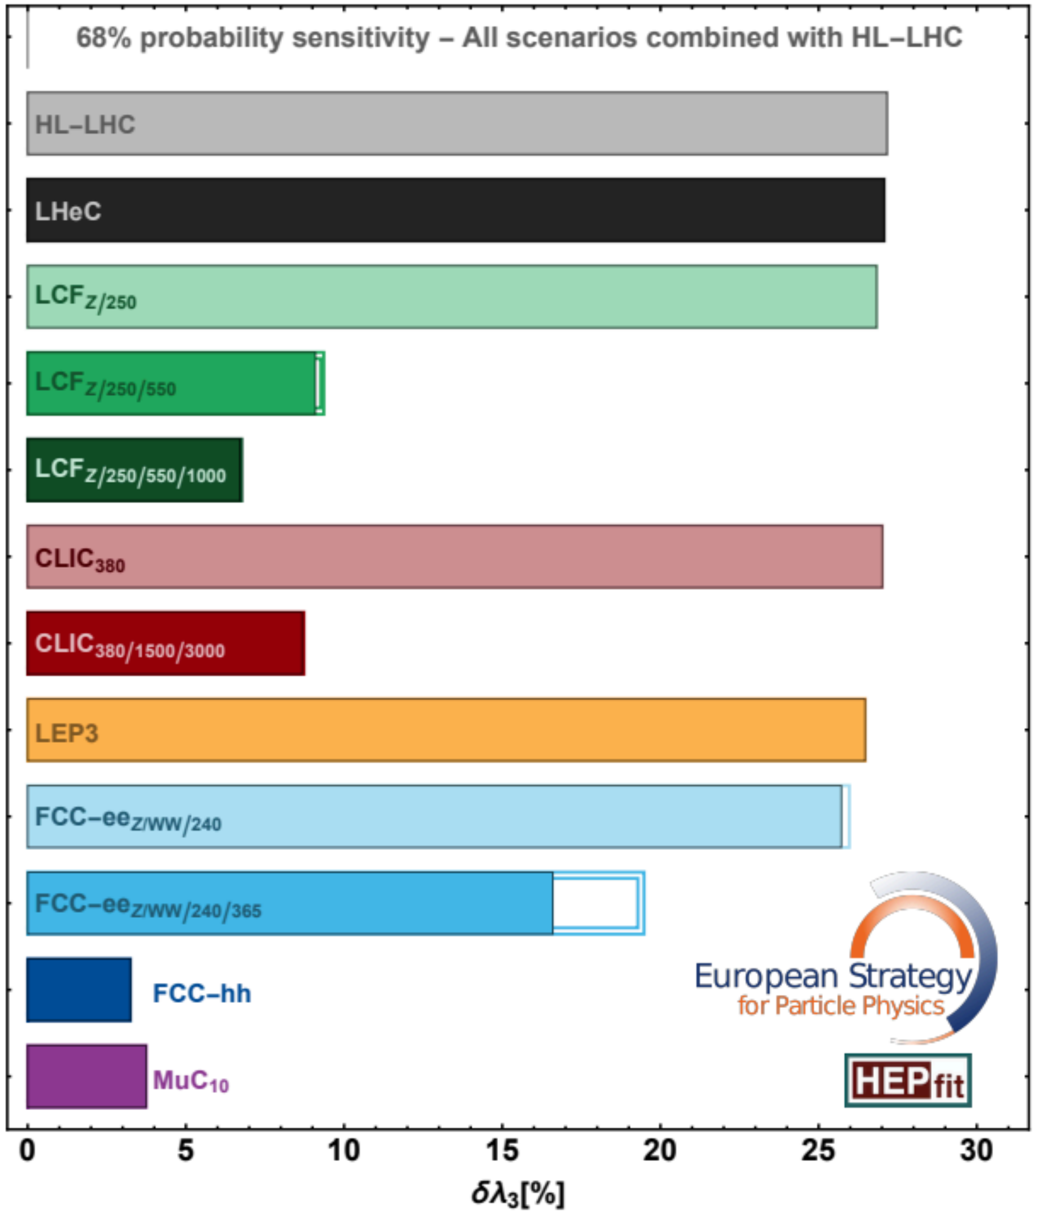

Main triumph (so far) of AI for particle physics? Enabling the discovery of the Higgs self-coupling at the LHC!

SMEFT interpretation

- Model-agnostic strategy to **interpret constraints on** (heavy) **BSM dynamics** affecting (very) different sectors & energy regimes
- A **single SMEFT analysis** automatically constrains all UV models matched to it

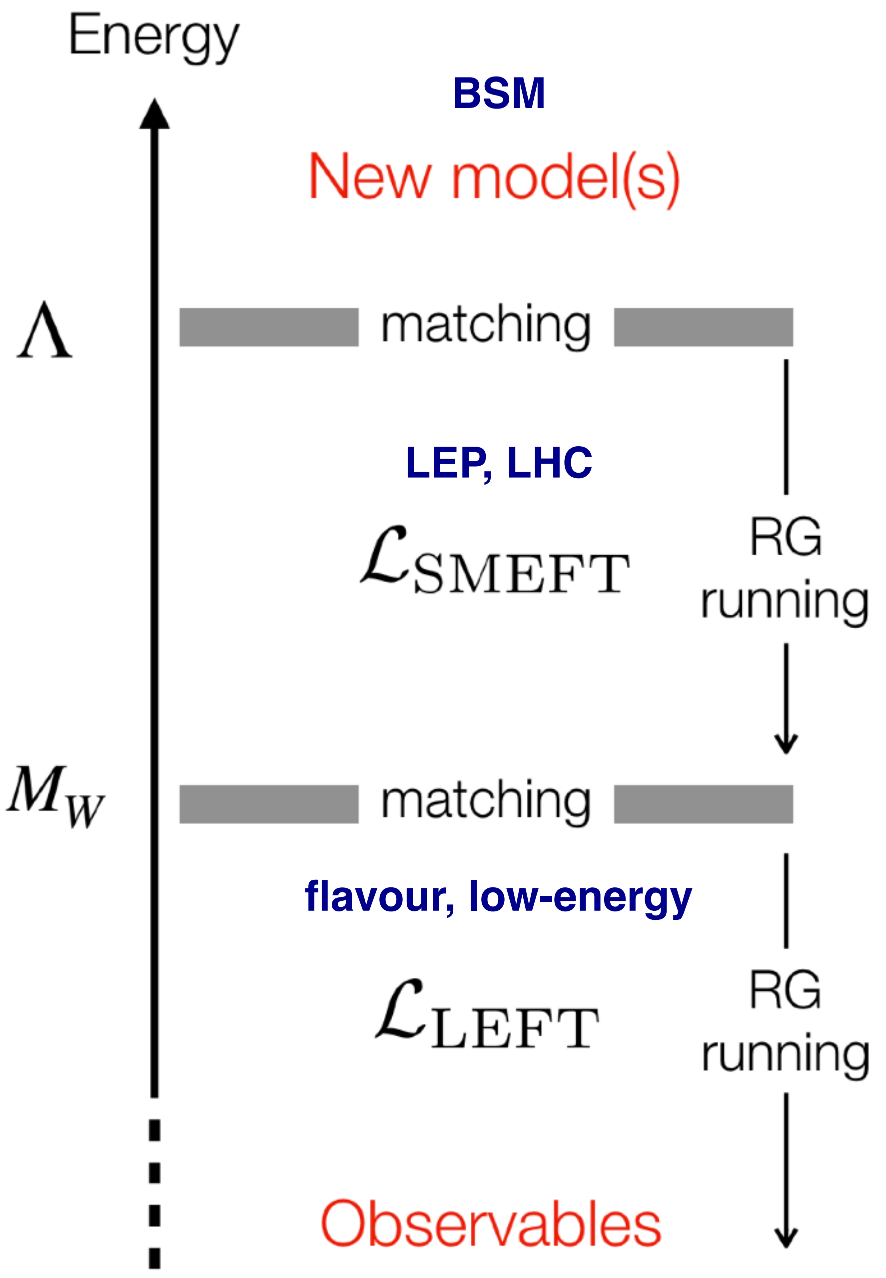

Energy

BSM

New model(s)

Λ matching

LEP, LHC

$\mathcal{L}_{\text{SMEFT}}$

RG running

M_W matching

flavour, low-energy

$\mathcal{L}_{\text{LEFT}}$

RG running

Observables

J. Fuentes-Martin, SM@LHC2024

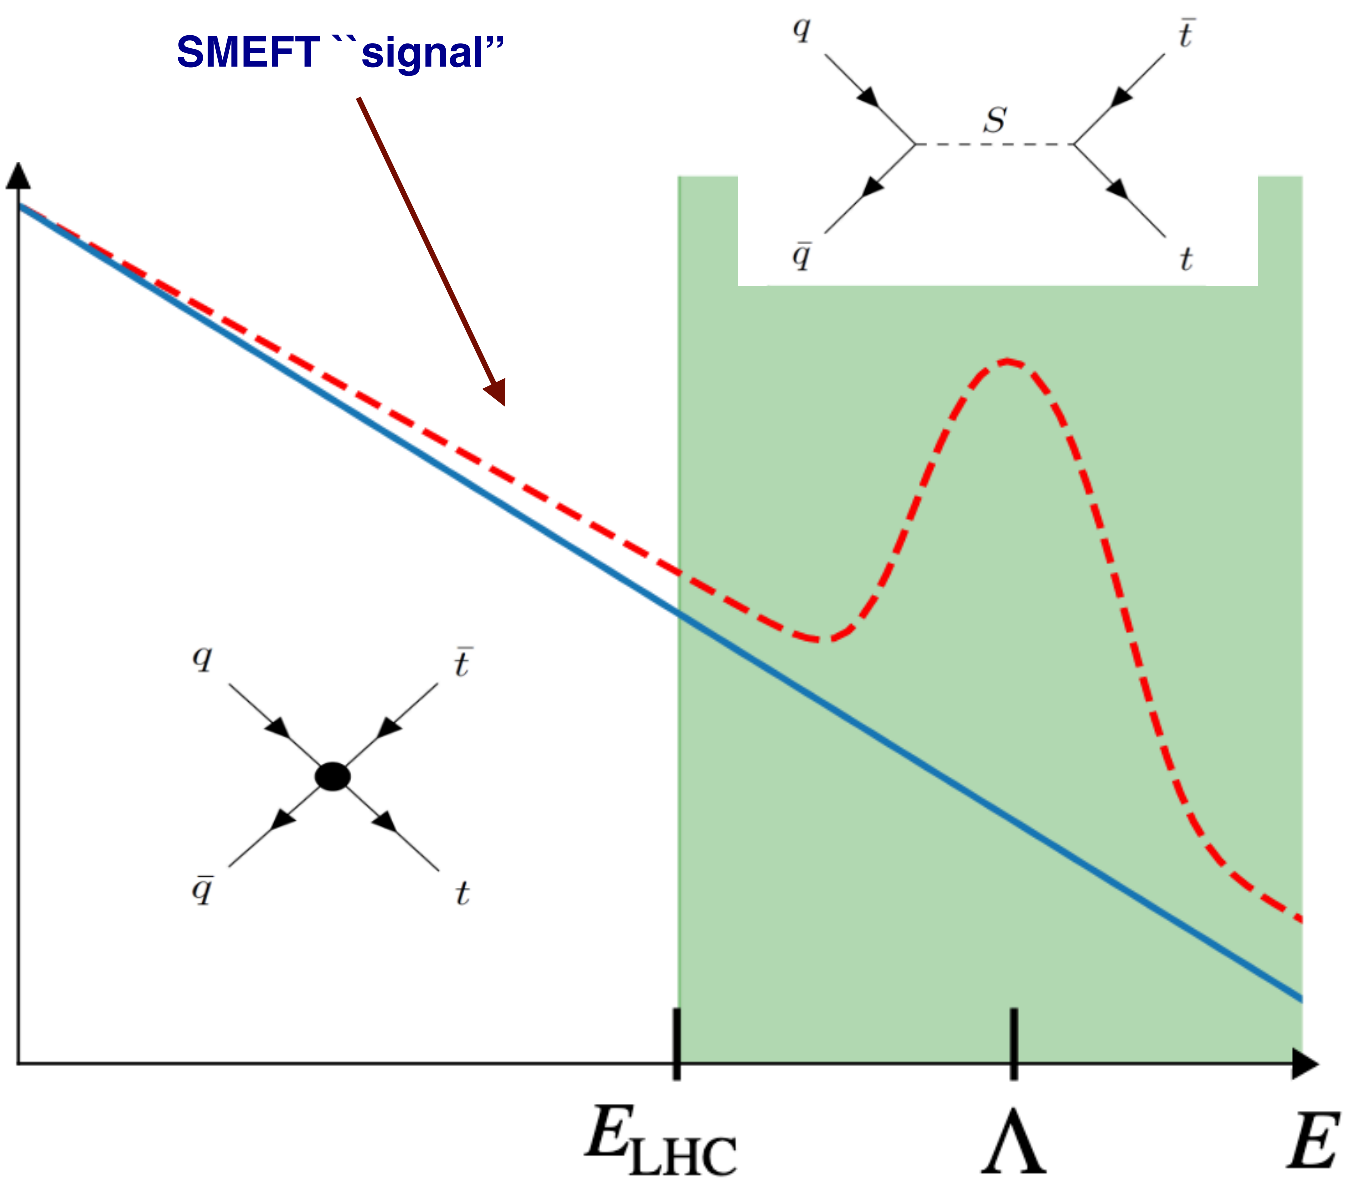

SMEFT ``signal''

SMEFT interpretation

- Specify the **particle (matter) content**: three generations of quarks and leptons
- Indicate **gauge** (local) symmetries and their eventual breaking mechanisms
- Impose **Lorentz** invariance and other global symmetries
- Ensure that predictions are valid **up to a cutoff scale Λ** (say Λ=10 TeV)

Wilson coefficients

$$\mathscr{L}_{\text{SMEFT}} = \mathscr{L}_{\text{SM}} + \sum_i \frac{c_i^{(5)}}{\Lambda}\mathcal{O}_i^{(d=5)} + \sum_j \frac{c_j^{(6)}}{\Lambda^2}\mathcal{O}_j^{(d=6)} + \sum_k \frac{c_j^{(7)}}{\Lambda^3}\mathcal{O}_k^{(d=7)}\ldots$$

cutoff (BSM) scale

all possible operators consistent
with these requirements

- The **dim-6 SMEFT** is the most relevant for LHC data (with exceptions *e.g.* VBS)

$$\sigma_{\text{SMEFT}}\left(c_j^{(6)}\right) = \sigma_{\text{SM}} + \sum_j \frac{c_j^{(6)}}{\Lambda^2}\sigma_{\text{EFT,j}} + \sum_{j,k} \frac{c_j^{(6)}c_k^{(6)}}{\Lambda^4}\sigma_{\text{EFT,jk}}$$

constrain from data

SMEFT interpretation

- Global SMEFT analysis of (HL-)LHC and future collider projections
- Study impact RGEs, **NLO QCD/EW**, higher order EFT, UV matching, theory uncertainties

$$\delta\kappa_3 = -\frac{2v^4}{m_h^2}\frac{c_\varphi}{\Lambda^2} + \frac{3v^2}{\Lambda^2}\left(c_{\varphi\Box} - \frac{1}{4}c_{\varphi D}\right)$$

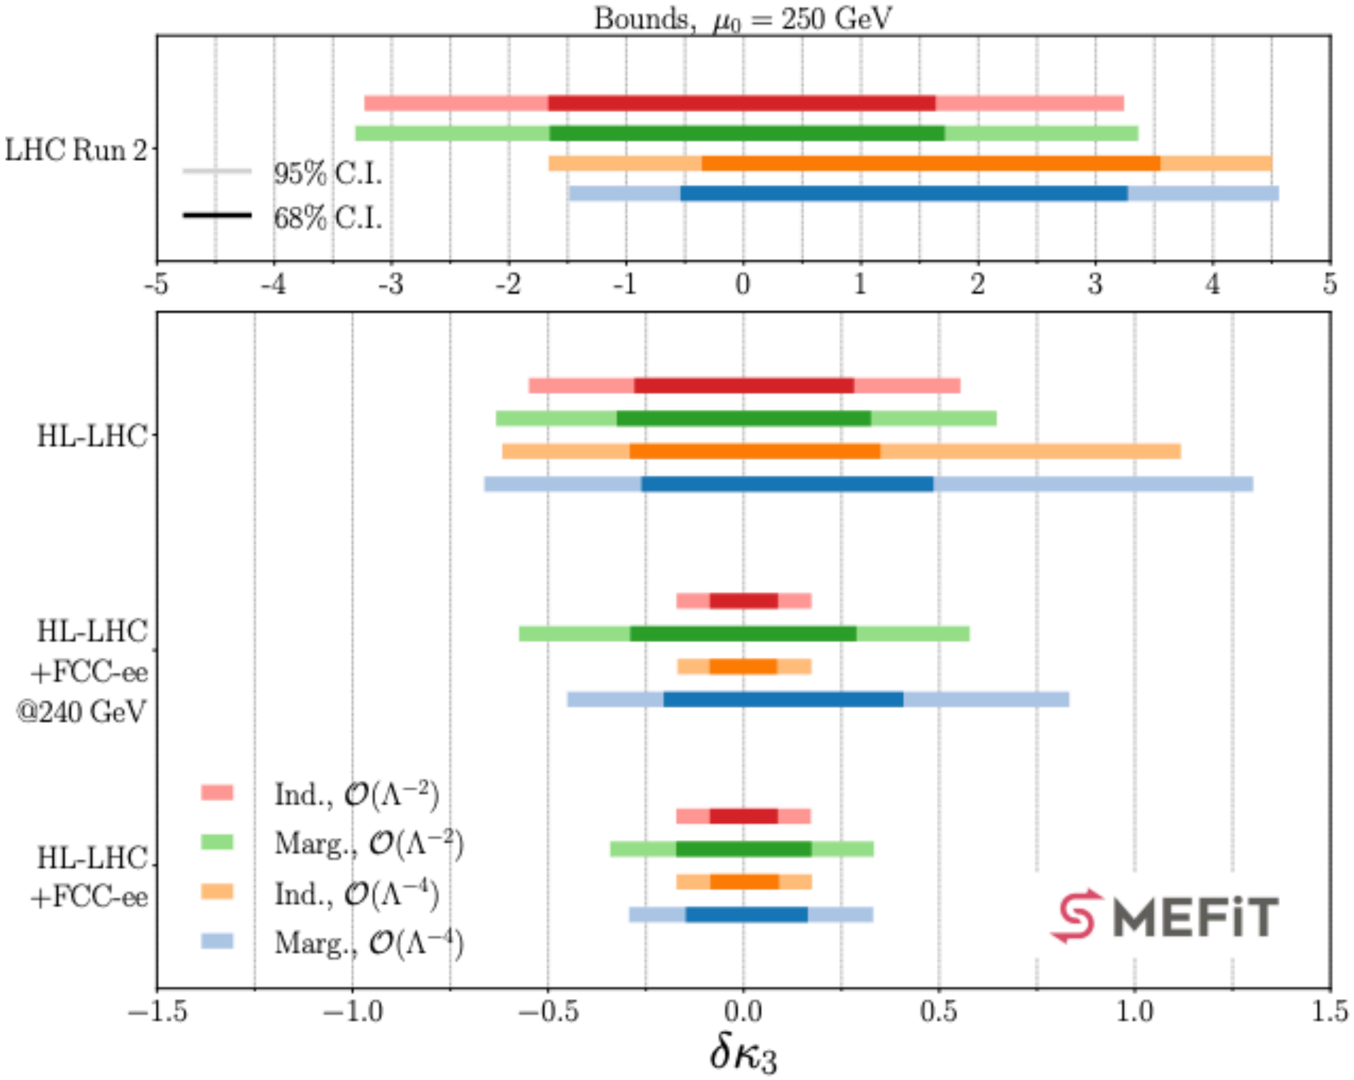

Input Dataset	EFT	$\delta\kappa_3$ (68% C.I.)
LHC Run 2	Linear	$[-1.68, 1.68]$
	Quad.	$[-0.54, 3.27]$
HL-LHC	Linear	$[-0.32, 0.32]$
	Quad.	$[-0.26, 0.49]$
HL-LHC & FCC-ee(240)	Linear	$[-0.29, 0.29]$
	Quad.	$[-0.20, 0.41]$
HL-LHC & FCC-ee	Linear	$[-0.17, 0.17]$
	Quad.	$[-0.15, 0.16]$

Large differences between one-parameter and marginalised fits: **global interpretation** unavoidable

FCC-ee needs top run to be competitive

ter Hoeve, Mantani, JR, Rossia, Vryonidiu (2025)

Agentic AI for HEP

Agentic AI for HEP

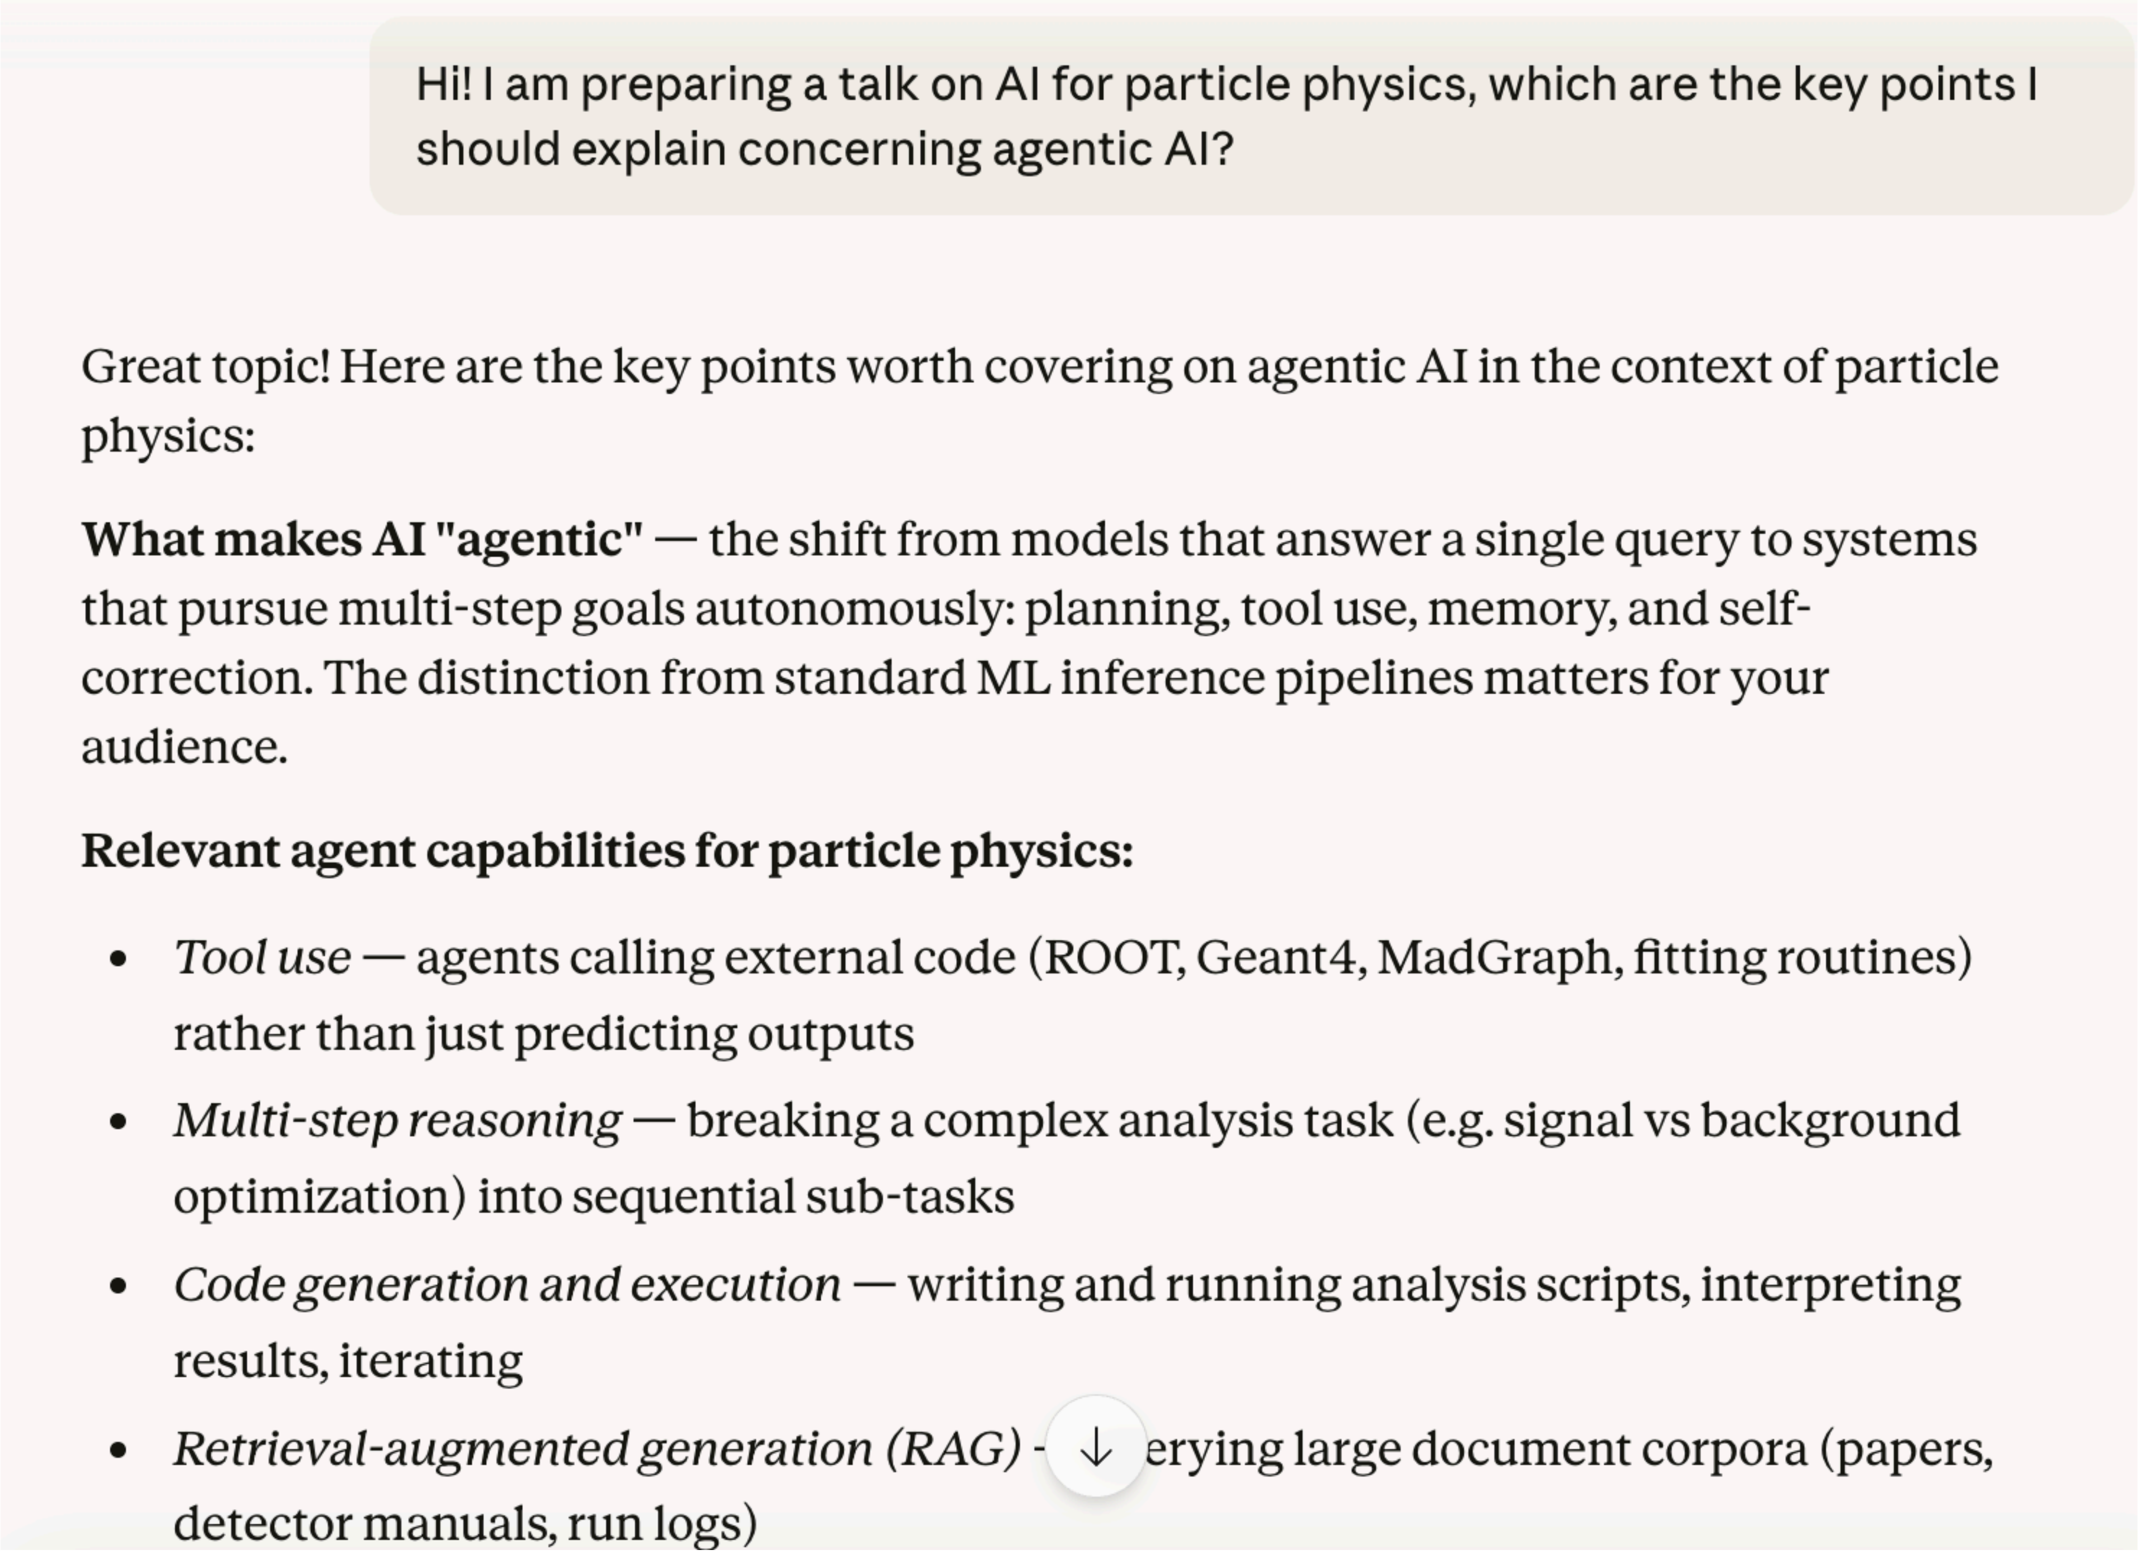

Hi! I am preparing a talk on AI for particle physics, which are the key points I should explain concerning agentic AI?

Great topic! Here are the key points worth covering on agentic AI in the context of particle physics:

What makes AI "agentic" — the shift from models that answer a single query to systems that pursue multi-step goals autonomously: planning, tool use, memory, and self-correction. The distinction from standard ML inference pipelines matters for your audience.

Relevant agent capabilities for particle physics:

- *Tool use* — agents calling external code (ROOT, Geant4, MadGraph, fitting routines) rather than just predicting outputs
- *Multi-step reasoning* — breaking a complex analysis task (e.g. signal vs background optimization) into sequential sub-tasks
- *Code generation and execution* — writing and running analysis scripts, interpreting results, iterating
- *Retrieval-augmented generation (RAG)* - erying large document corpora (papers, detector manuals, run logs)

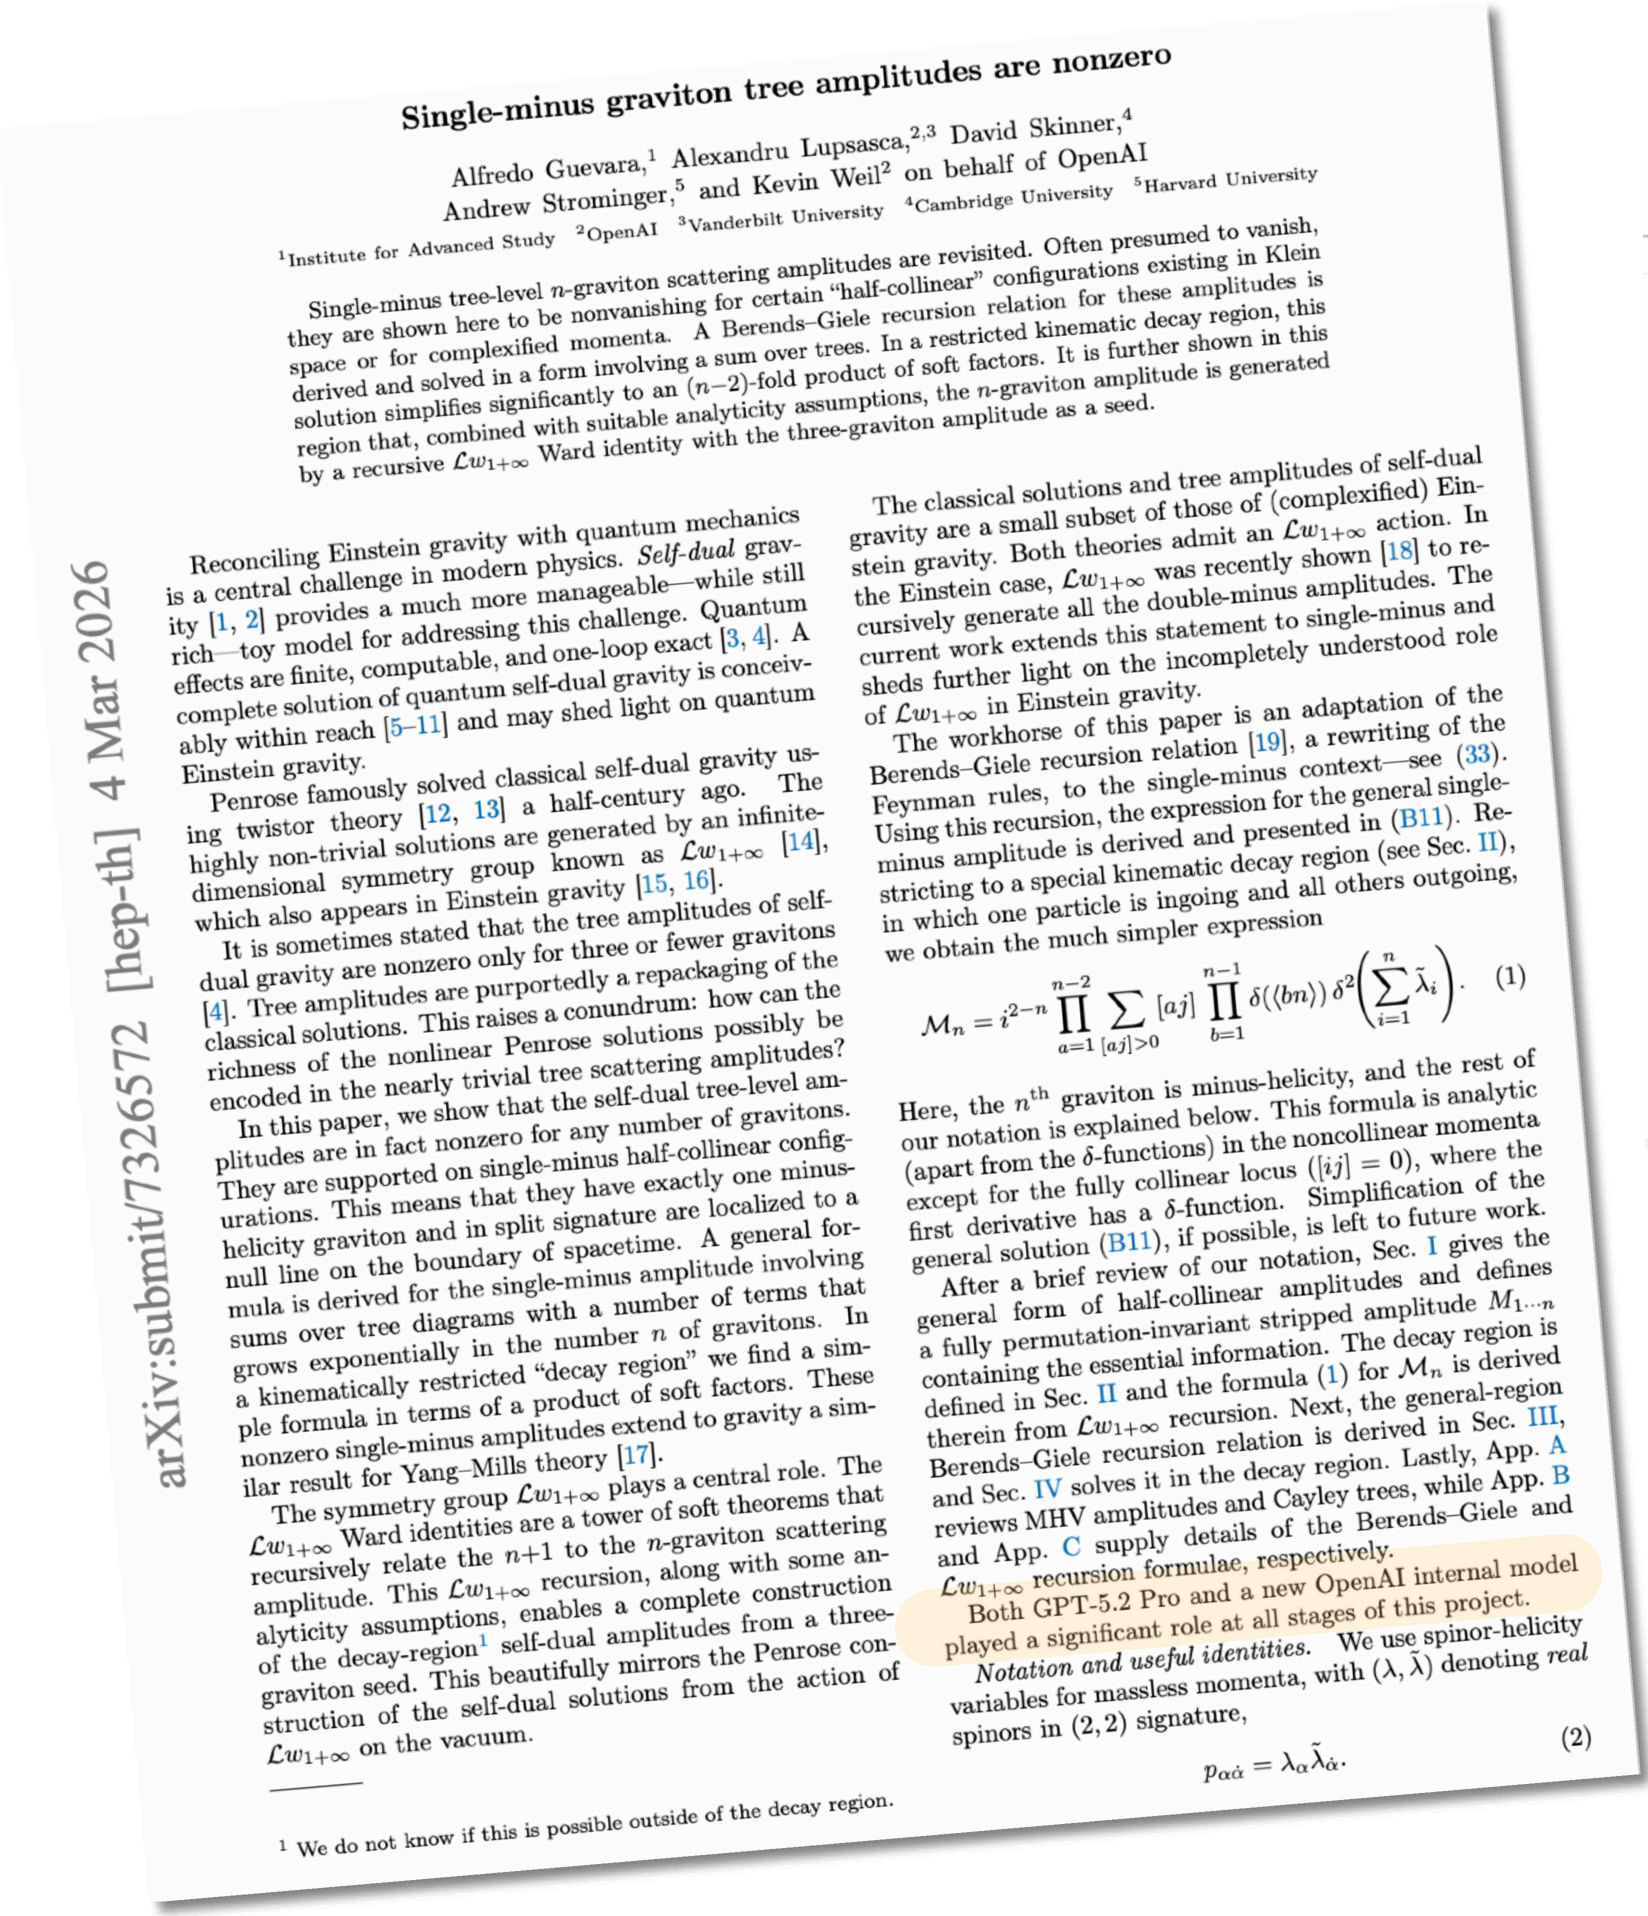

Single-minus graviton tree amplitudes are nonzero

Alfredo Guevara,[1] Alexandru Lupsasca,[2,3] David Skinner,[4]
Andrew Strominger,[5] and Kevin Weil[2] on behalf of OpenAI
[1]Institute for Advanced Study [2]OpenAI [3]Vanderbilt University [4]Cambridge University [5]Harvard University

Single-minus tree-level n-graviton scattering amplitudes are revisited. Often presumed to vanish, they are shown here to be nonvanishing for certain "half-collinear" configurations existing in Klein space or for complexified momenta. A Berends–Giele recursion relation for these amplitudes is derived and solved in a form involving a sum over trees. In a restricted kinematic decay region, this solution simplifies significantly to an $(n-2)$-fold product of soft factors. It is further shown in this region that, combined with suitable analyticity assumptions, the n-graviton amplitude is generated by a recursive $\mathcal{L}w_{1+\infty}$ Ward identity with the three-graviton amplitude as a seed.

Reconciling Einstein gravity with quantum mechanics is a central challenge in modern physics. *Self-dual* gravity [1, 2] provides a much more manageable—while still rich—toy model for addressing this challenge. Quantum effects are finite, computable, and one-loop exact [3, 4]. A complete solution of quantum self-dual gravity is conceivably within reach [5–11] and may shed light on quantum Einstein gravity.

Penrose famously solved classical self-dual gravity using twistor theory [12, 13] a half-century ago. The highly non-trivial solutions are generated by an infinite-dimensional symmetry group known as $\mathcal{L}w_{1+\infty}$ [14], which also appears in Einstein gravity [15, 16].

It is sometimes stated that the tree amplitudes of self-dual gravity are nonzero only for three or fewer gravitons [4]. Tree amplitudes are purportedly a repackaging of the classical solutions. This raises a conundrum: how can the richness of the nonlinear Penrose solutions possibly be encoded in the nearly trivial tree scattering amplitudes?

In this paper, we show that the self-dual tree-level amplitudes are in fact nonzero for any number of gravitons. They are supported on single-minus half-collinear configurations. This means that they have exactly one minus-helicity graviton and in split signature are localized to a null line on the boundary of spacetime. A general formula is derived for the single-minus amplitude involving sums over tree diagrams with a number of terms that grows exponentially in the number n of gravitons. In a kinematically restricted "decay region" we find a simple formula in terms of a product of soft factors. These nonzero single-minus amplitudes extend to gravity a similar result for Yang–Mills theory [17].

The symmetry group $\mathcal{L}w_{1+\infty}$ plays a central role. The $\mathcal{L}w_{1+\infty}$ Ward identities are a tower of soft theorems that recursively relate the $n{+}1$ to the n-graviton scattering amplitude. This $\mathcal{L}w_{1+\infty}$ recursion, along with some analyticity assumptions, enables a complete construction of the decay-region[1] self-dual amplitudes from a three-graviton seed. This beautifully mirrors the Penrose construction of the self-dual solutions from the action of $\mathcal{L}w_{1+\infty}$ on the vacuum.

The classical solutions and tree amplitudes of self-dual gravity are a small subset of those of (complexified) Einstein gravity. Both theories admit an $\mathcal{L}w_{1+\infty}$ action. In the Einstein case, $\mathcal{L}w_{1+\infty}$ was recently shown [18] to recursively generate all the double-minus amplitudes. The current work extends this statement to single-minus and sheds further light on the incompletely understood role of $\mathcal{L}w_{1+\infty}$ in Einstein gravity.

The workhorse of this paper is an adaptation of the Berends–Giele recursion relation [19], a rewriting of the Feynman rules, to the single-minus context—see (33). Using this recursion, the expression for the general single-minus amplitude is derived and presented in (B11). Restricting to a special kinematic decay region (see Sec. II), in which one particle is ingoing and all others outgoing, we obtain the much simpler expression

$$\mathcal{M}_n = i^{2-n}\prod_{a=1}^{n-2}\sum_{[aj]>0}[aj]\prod_{b=1}^{n-1}\delta(\langle bn\rangle)\,\delta^2\Bigg(\sum_{i=1}^{n}\tilde\lambda_i\Bigg). \qquad (1)$$

Here, the nth graviton is minus-helicity, and the rest of our notation is explained below. This formula is analytic (apart from the δ-functions) in the noncollinear momenta except for the fully collinear locus ($[ij]=0$), where the first derivative has a δ-function. Simplification of the general solution (B11), if possible, is left to future work.

After a brief review of our notation, Sec. I gives the general form of half-collinear amplitudes and defines a fully permutation-invariant stripped amplitude $M_{1\cdots n}$ containing the essential information. The decay region is defined in Sec. II and the formula (1) for $\mathcal{M}_n$ is derived therein from $\mathcal{L}w_{1+\infty}$ recursion. Next, the general-region Berends–Giele recursion relation is derived in Sec. III, and Sec. IV solves it in the decay region. Lastly, App. A reviews MHV amplitudes and Cayley trees, while App. B and App. C supply details of the Berends–Giele and $\mathcal{L}w_{1+\infty}$ recursion formulae, respectively. Both GPT-5.2 Pro and a new OpenAI internal model played a significant role at all stages of this project.

Notation and useful identities. We use spinor-helicity variables for massless momenta, with $(\lambda,\tilde\lambda)$ denoting *real* spinors in $(2,2)$ signature,

$$p_{\alpha\dot\alpha} = \lambda_\alpha\tilde\lambda_{\dot\alpha}. \qquad (2)$$

---

[1] We do not know if this is possible outside of the decay region.

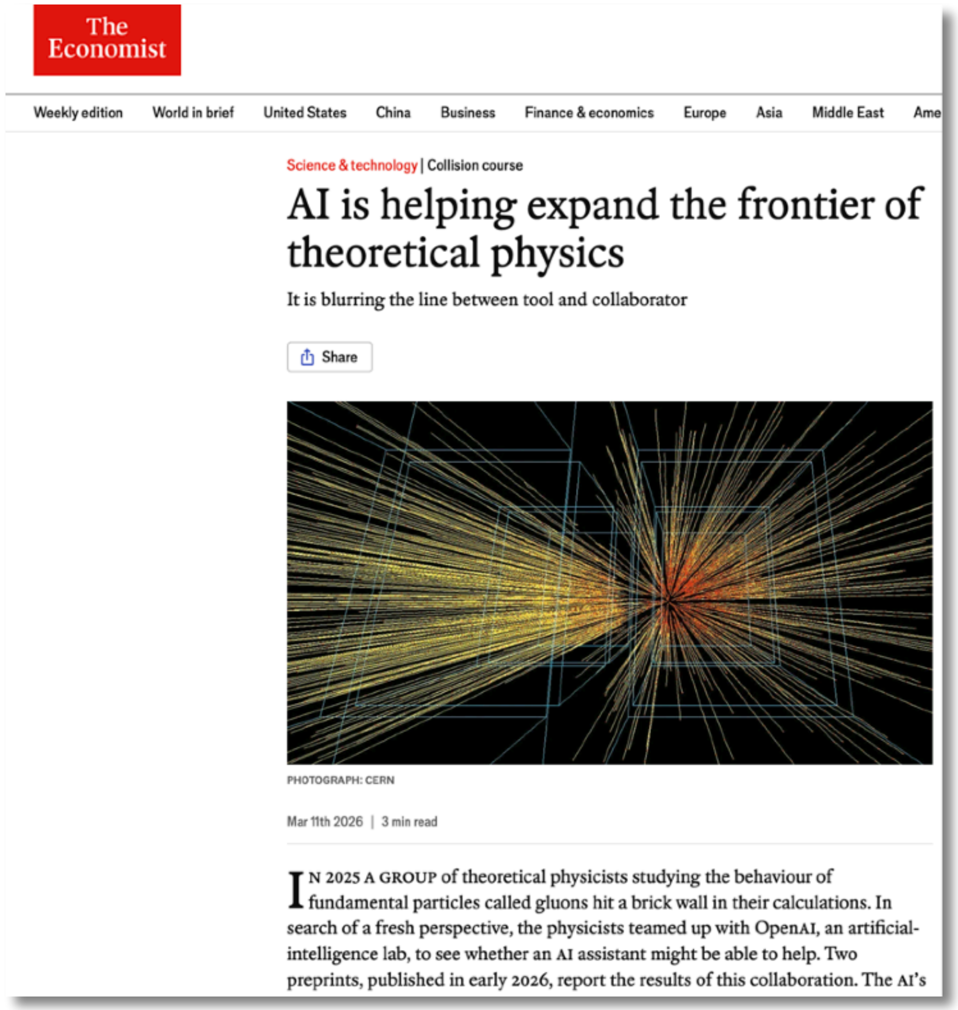

Science & technology | Collision course

AI is helping expand the frontier of theoretical physics

It is blurring the line between tool and collaborator

PHOTOGRAPH: CERN

Mar 11th 2026 | 3 min read

IN 2025 A GROUP of theoretical physicists studying the behaviour of fundamental particles called gluons hit a brick wall in their calculations. In search of a fresh perspective, the physicists teamed up with OpenAI, an artificial-intelligence lab, to see whether an AI assistant might be able to help. Two preprints, published in early 2026, report the results of this collaboration. The AI's

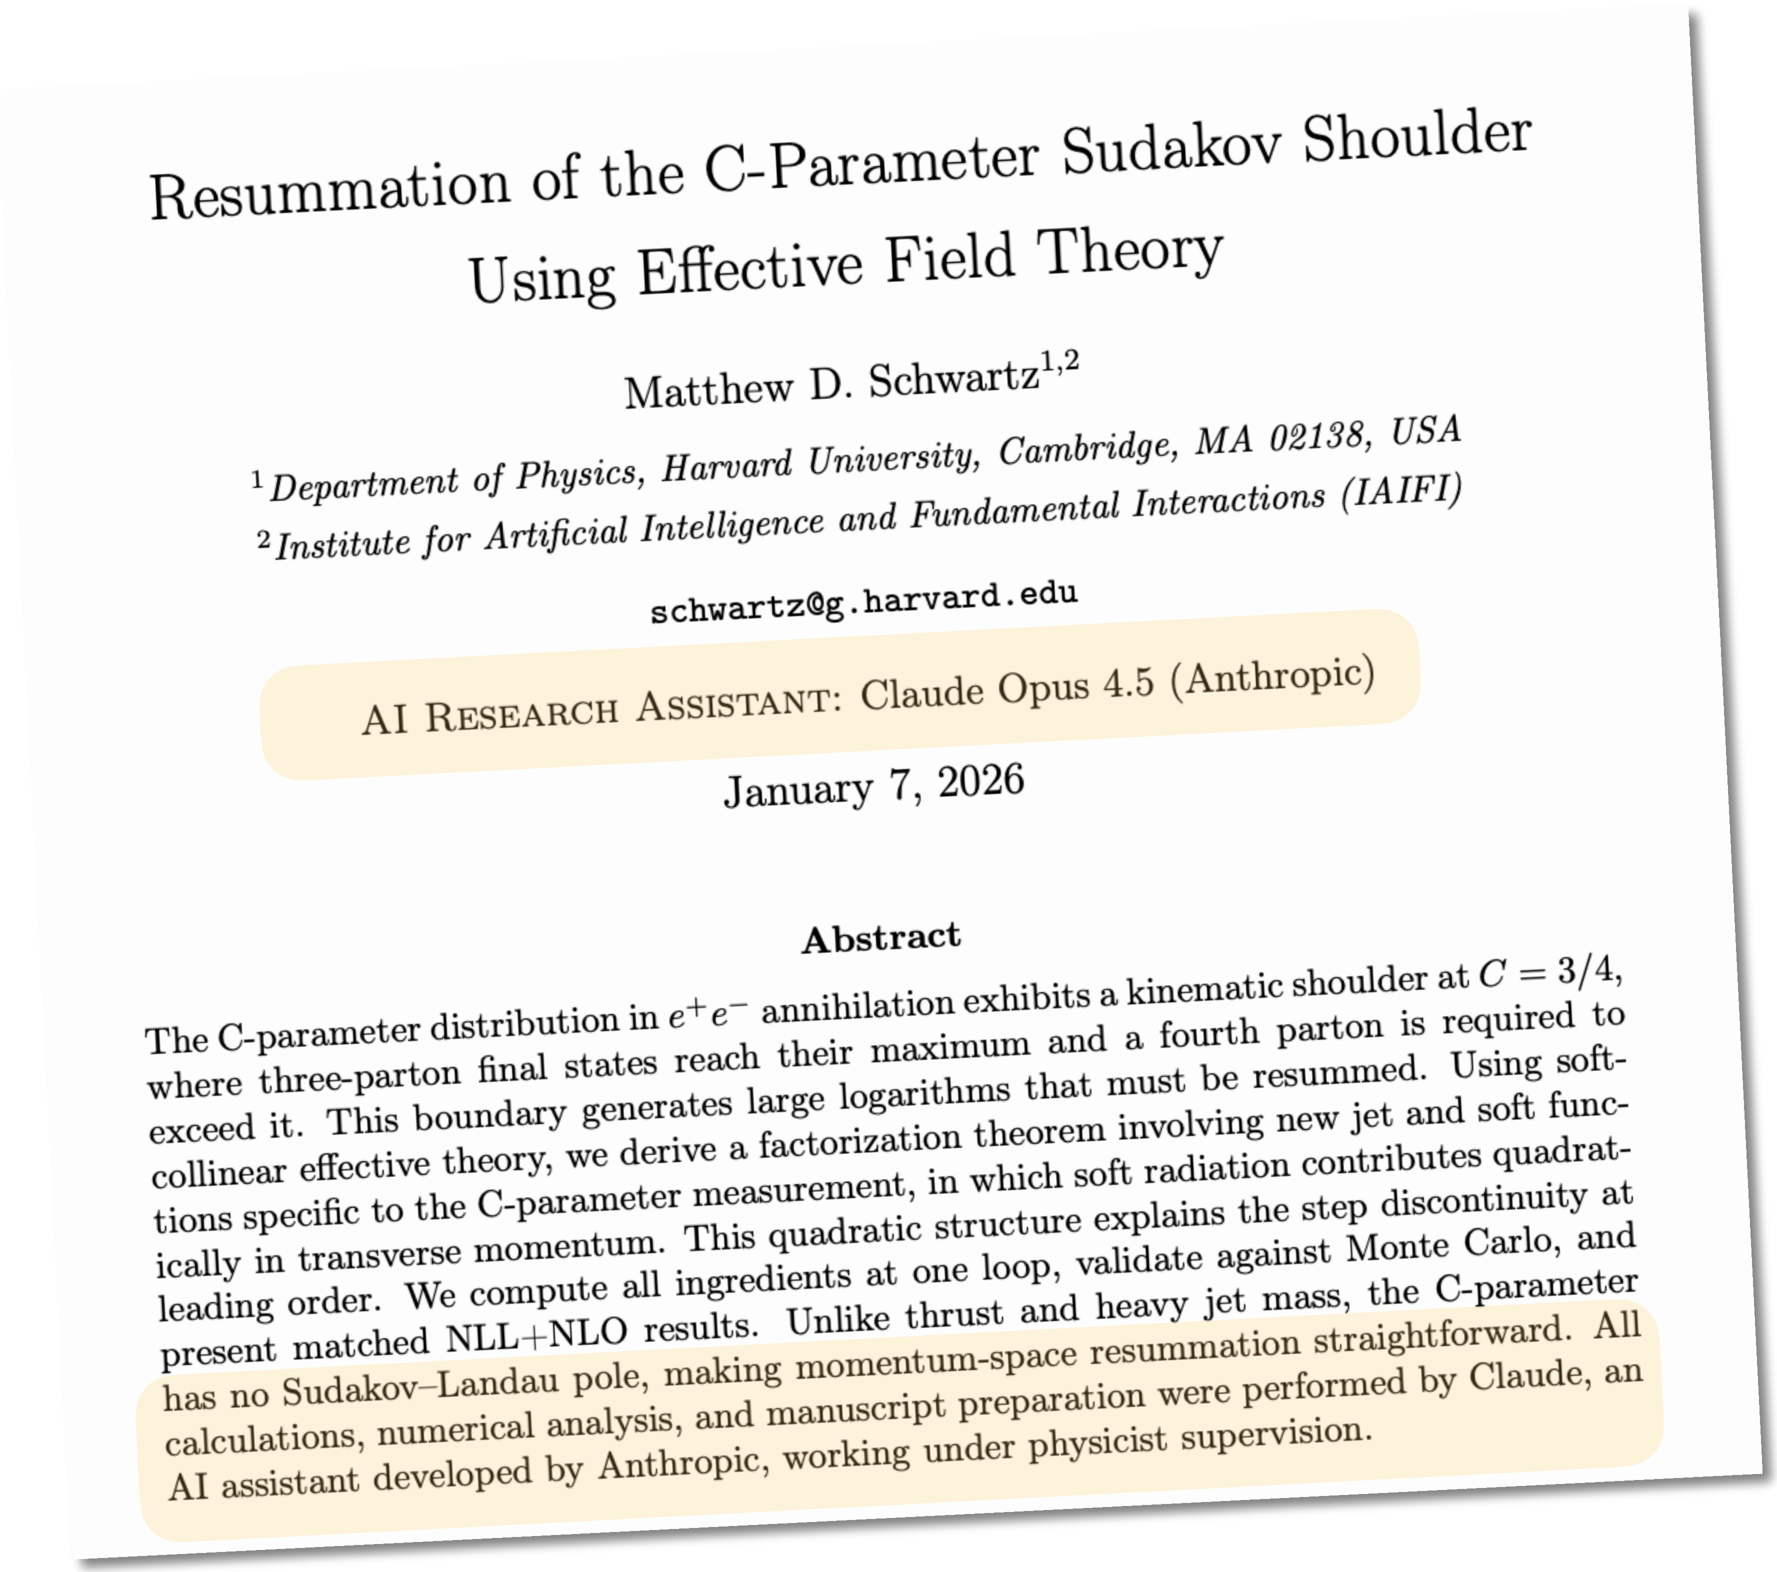

Resummation of the C-Parameter Sudakov Shoulder Using Effective Field Theory

Matthew D. Schwartz[1,2]

[1] *Department of Physics, Harvard University, Cambridge, MA 02138, USA*

[2] *Institute for Artificial Intelligence and Fundamental Interactions (IAIFI)*

`schwartz@g.harvard.edu`

AI Research Assistant: Claude Opus 4.5 (Anthropic)

January 7, 2026

Abstract

The C-parameter distribution in e^+e^- annihilation exhibits a kinematic shoulder at $C = 3/4$, where three-parton final states reach their maximum and a fourth parton is required to exceed it. This boundary generates large logarithms that must be resummed. Using soft-collinear effective theory, we derive a factorization theorem involving new jet and soft functions specific to the C-parameter measurement, in which soft radiation contributes quadratically in transverse momentum. This quadratic structure explains the step discontinuity at leading order. We compute all ingredients at one loop, validate against Monte Carlo, and present matched NLL+NLO results. Unlike thrust and heavy jet mass, the C-parameter has no Sudakov–Landau pole, making momentum-space resummation straightforward. All calculations, numerical analysis, and manuscript preparation were performed by Claude, an AI assistant developed by Anthropic, working under physicist supervision.

Productivity-enhancing tools or replacing physicists?

MadAgents

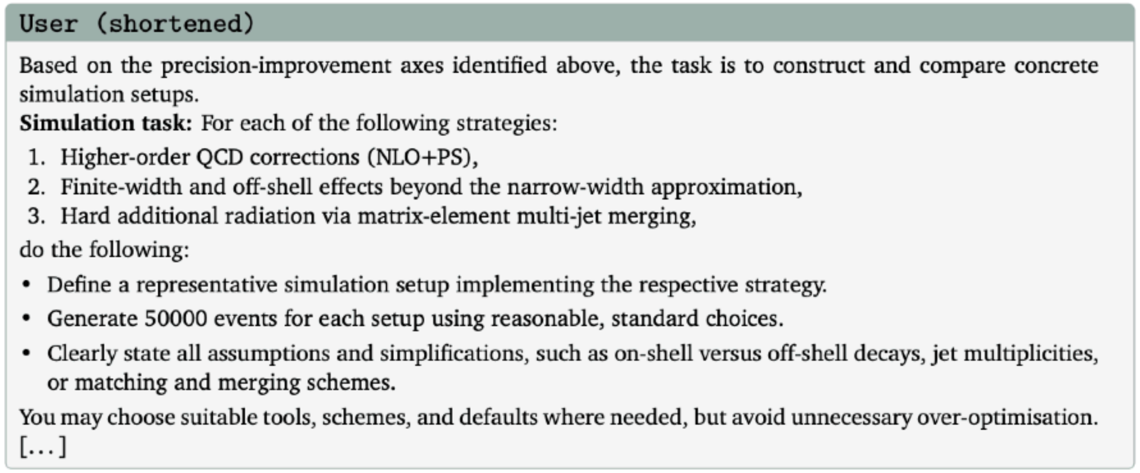

`User (shortened)`

Based on the precision-improvement axes identified above, the task is to construct and compare concrete simulation setups.

Simulation task: For each of the following strategies:

1. Higher-order QCD corrections (NLO+PS),
2. Finite-width and off-shell effects beyond the narrow-width approximation,
3. Hard additional radiation via matrix-element multi-jet merging,

do the following:

- Define a representative simulation setup implementing the respective strategy.
- Generate 50000 events for each setup using reasonable, standard choices.
- Clearly state all assumptions and simplifications, such as on-shell versus off-shell decays, jet multiplicities, or matching and merging schemes.

You may choose suitable tools, schemes, and defaults where needed, but avoid unnecessary over-optimisation.
[. . .]

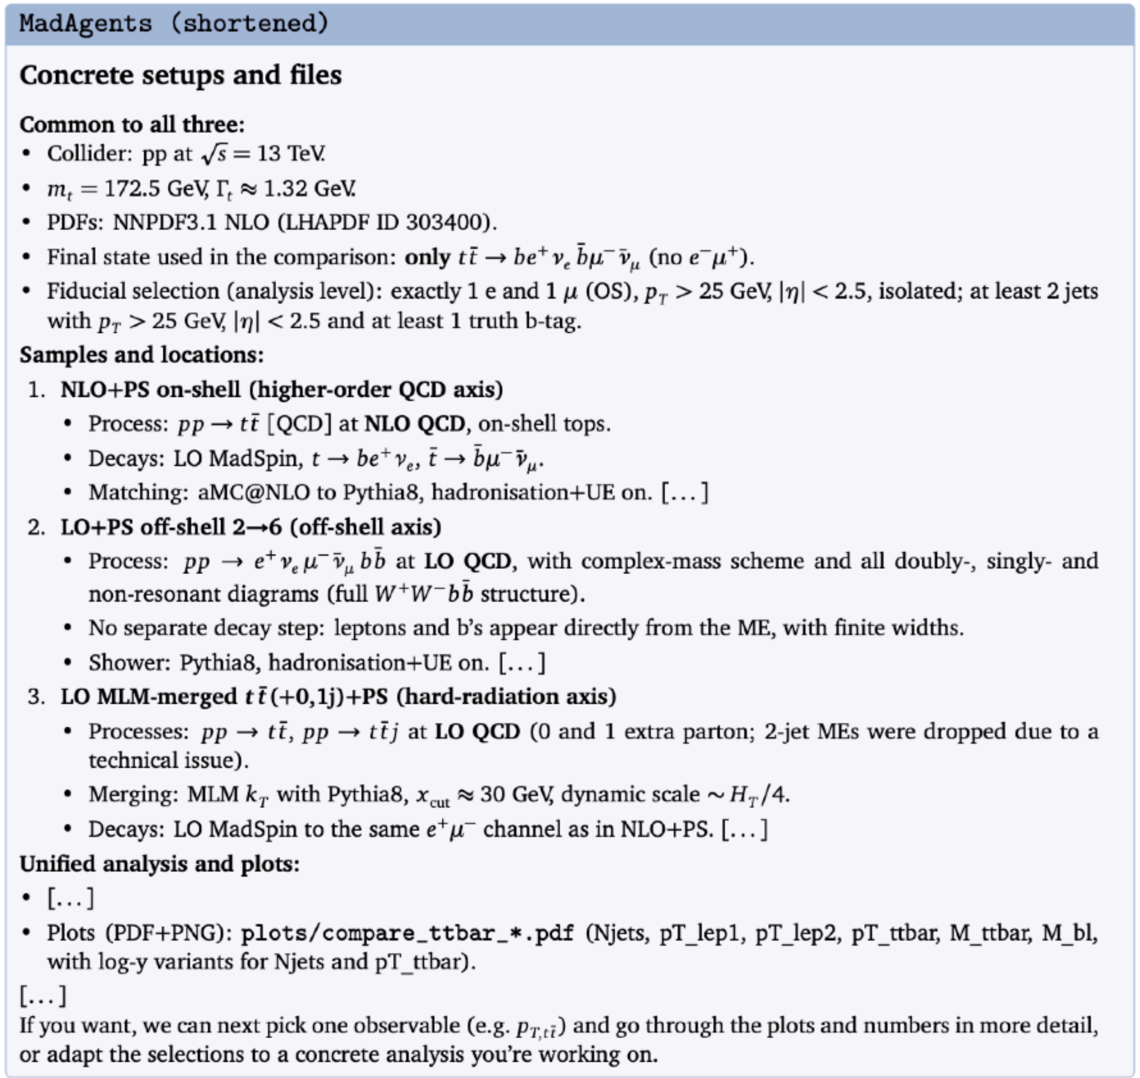

`MadAgents (shortened)`

Concrete setups and files

Common to all three:

- Collider: pp at $\sqrt{s} = 13$ TeV.
- $m_t = 172.5$ GeV, $\Gamma_t \approx 1.32$ GeV.
- PDFs: NNPDF3.1 NLO (LHAPDF ID 303400).
- Final state used in the comparison: **only** $t\bar{t} \to be^+\nu_e\bar{b}\mu^-\bar{\nu}_\mu$ (no $e^-\mu^+$).
- Fiducial selection (analysis level): exactly 1 e and 1 μ (OS), $p_T > 25$ GeV, $|\eta| < 2.5$, isolated; at least 2 jets with $p_T > 25$ GeV, $|\eta| < 2.5$ and at least 1 truth b-tag.

Samples and locations:

1. **NLO+PS on-shell (higher-order QCD axis)**
   - Process: $pp \to t\bar{t}$ [QCD] at **NLO QCD**, on-shell tops.
   - Decays: LO MadSpin, $t \to be^+\nu_e$, $\bar{t} \to \bar{b}\mu^-\bar{\nu}_\mu$.
   - Matching: aMC@NLO to Pythia8, hadronisation+UE on. [. . .]
2. **LO+PS off-shell 2→6 (off-shell axis)**
   - Process: $pp \to e^+\nu_e\mu^-\bar{\nu}_\mu b\bar{b}$ at **LO QCD**, with complex-mass scheme and all doubly-, singly- and non-resonant diagrams (full $W^+W^-b\bar{b}$ structure).
   - No separate decay step: leptons and b's appear directly from the ME, with finite widths.
   - Shower: Pythia8, hadronisation+UE on. [. . .]
3. **LO MLM-merged $t\bar{t}$(+0,1j)+PS (hard-radiation axis)**
   - Processes: $pp \to t\bar{t}$, $pp \to t\bar{t}j$ at **LO QCD** (0 and 1 extra parton; 2-jet MEs were dropped due to a technical issue).
   - Merging: MLM k_T with Pythia8, $x_{\text{cut}} \approx 30$ GeV, dynamic scale $\sim H_T/4$.
   - Decays: LO MadSpin to the same $e^+\mu^-$ channel as in NLO+PS. [. . .]

Unified analysis and plots:

- [. . .]
- Plots (PDF+PNG): **`plots/compare_ttbar_*.pdf`** (Njets, pT_lep1, pT_lep2, pT_ttbar, M_ttbar, M_bl, with log-y variants for Njets and pT_ttbar).

[. . .]
If you want, we can next pick one observable (e.g. $p_{T,t\bar{t}}$) and go through the plots and numbers in more detail, or adapt the selections to a concrete analysis you're working on.

- LHC simulations entirely prompt-driven
- Input the .pdf of a publication, produce **automatically the associated madgraph simulations**
- Assist user with result interpretation, though extensive literature searches

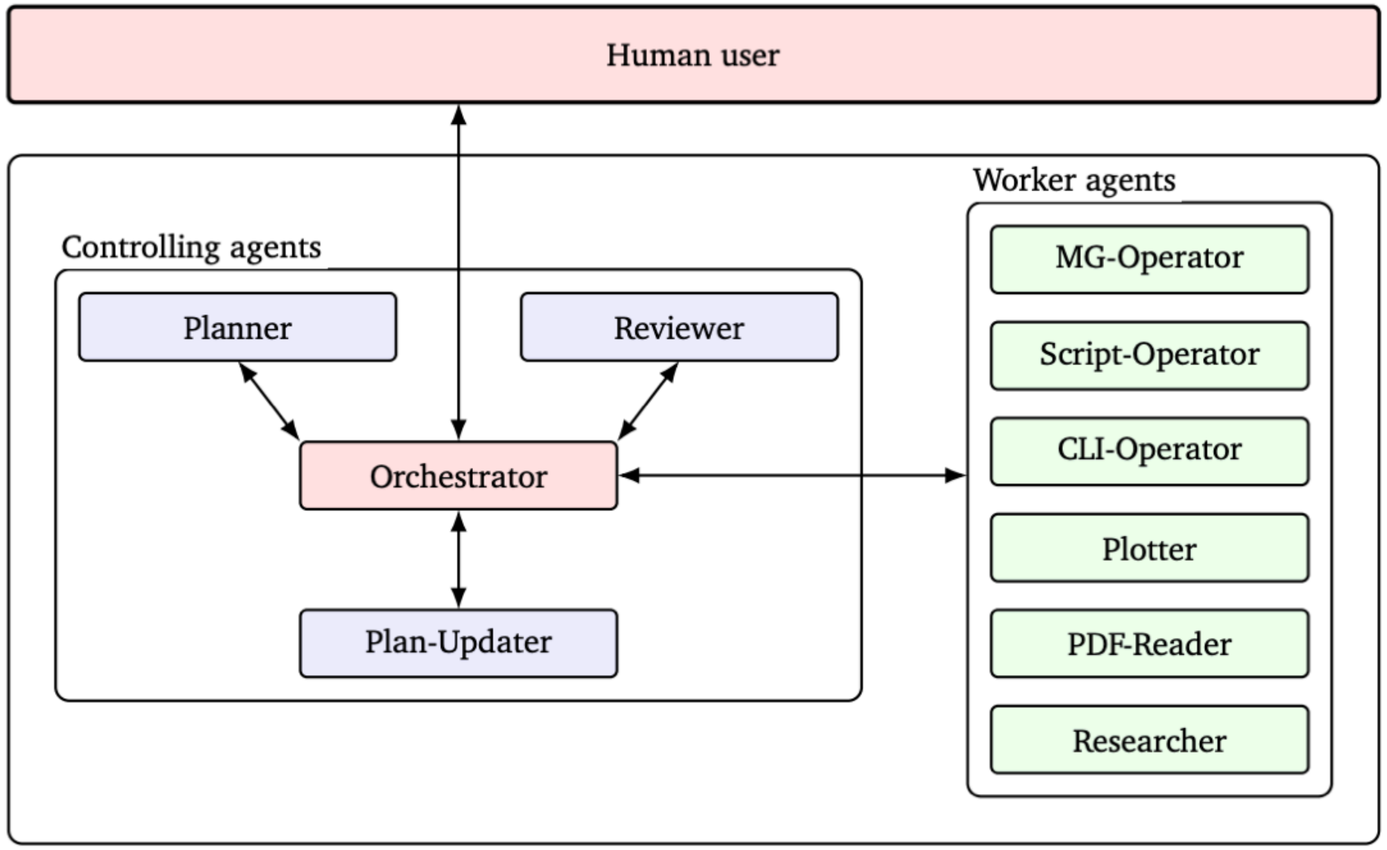

Human user

Worker agents

Controlling agents

Planner

Reviewer

Orchestrator

Plan-Updater

MG-Operator

Script-Operator

CLI-Operator

Plotter

PDF-Reader

Researcher

Autonomous Agents for Experimental HEP

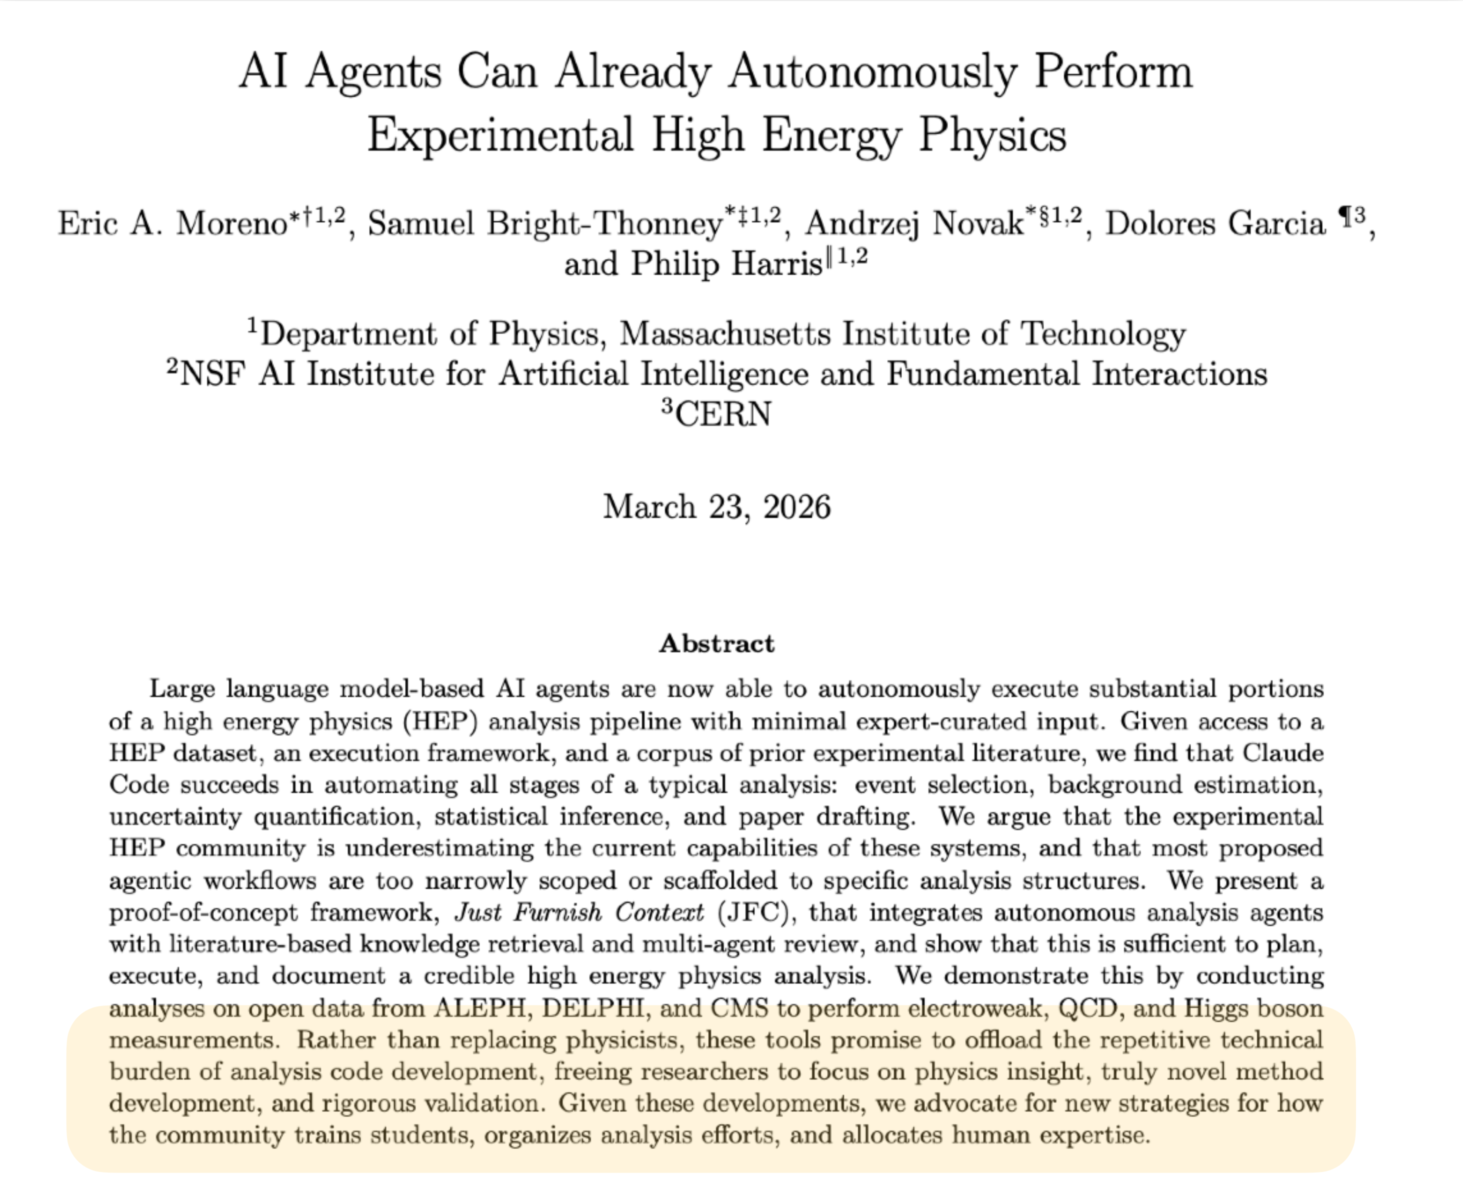

AI Agents Can Already Autonomously Perform Experimental High Energy Physics

Eric A. Moreno*†1,2, Samuel Bright-Thonney*‡1,2, Andrzej Novak*§1,2, Dolores Garcia ¶3, and Philip Harris‖1,2

1Department of Physics, Massachusetts Institute of Technology
2NSF AI Institute for Artificial Intelligence and Fundamental Interactions
3CERN

March 23, 2026

Abstract

Large language model-based AI agents are now able to autonomously execute substantial portions of a high energy physics (HEP) analysis pipeline with minimal expert-curated input. Given access to a HEP dataset, an execution framework, and a corpus of prior experimental literature, we find that Claude Code succeeds in automating all stages of a typical analysis: event selection, background estimation, uncertainty quantification, statistical inference, and paper drafting. We argue that the experimental HEP community is underestimating the current capabilities of these systems, and that most proposed agentic workflows are too narrowly scoped or scaffolded to specific analysis structures. We present a proof-of-concept framework, *Just Furnish Context* (JFC), that integrates autonomous analysis agents with literature-based knowledge retrieval and multi-agent review, and show that this is sufficient to plan, execute, and document a credible high energy physics analysis. We demonstrate this by conducting analyses on open data from ALEPH, DELPHI, and CMS to perform electroweak, QCD, and Higgs boson measurements. Rather than replacing physicists, these tools promise to offload the repetitive technical burden of analysis code development, freeing researchers to focus on physics insight, truly novel method development, and rigorous validation. Given these developments, we advocate for new strategies for how the community trains students, organizes analysis efforts, and allocates human expertise.

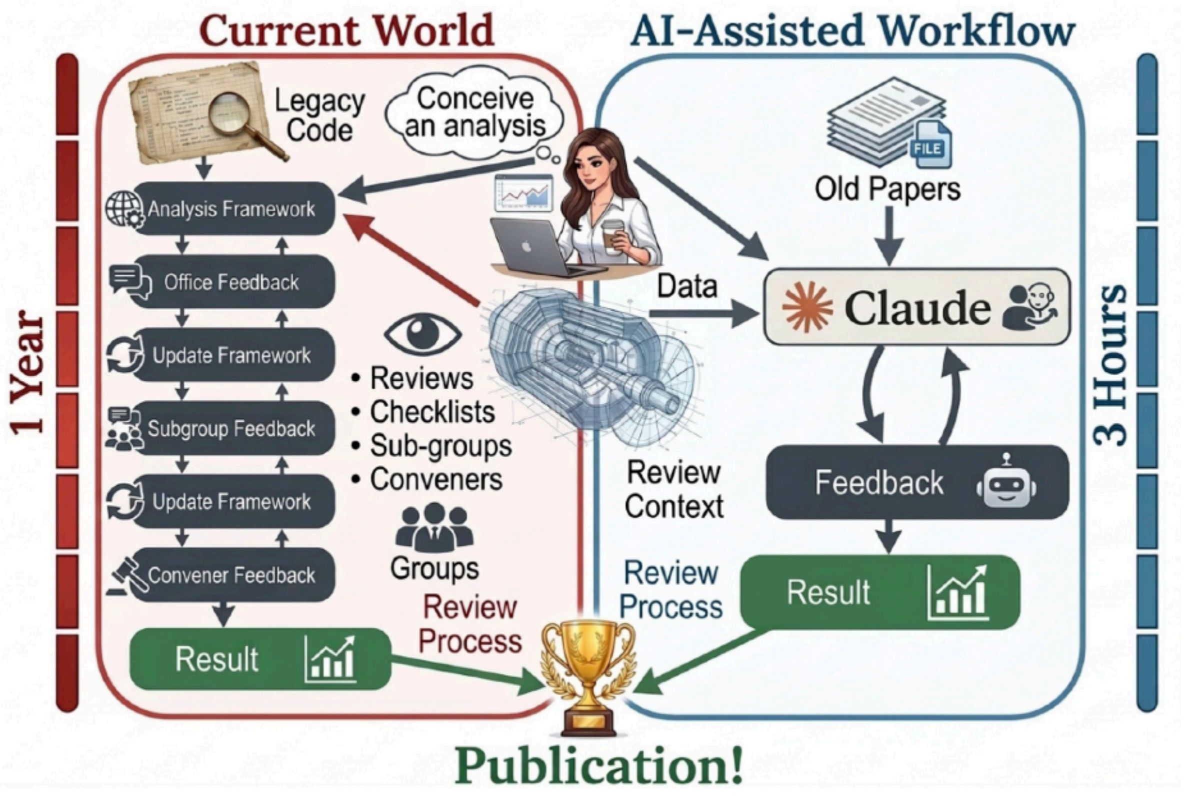

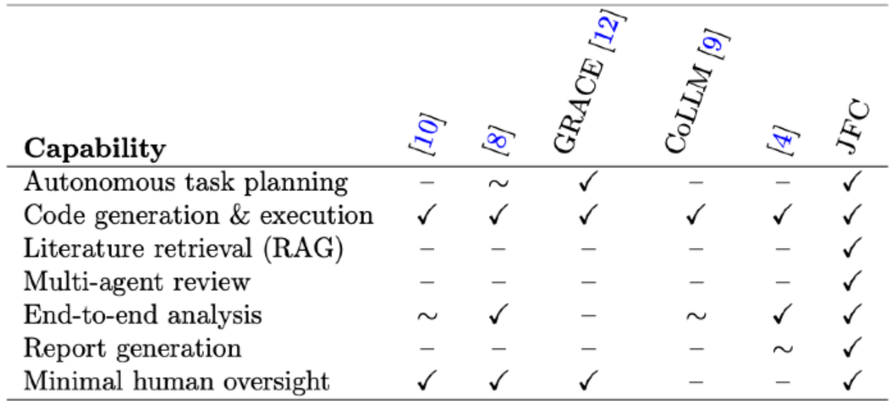

Capability	[10]	[8]	GRACE [12]	CoLLM [9]	[4]	JFC
Autonomous task planning	–	∼	✓	–	–	✓
Code generation & execution	✓	✓	✓	✓	✓	✓
Literature retrieval (RAG)	–	–	–	–	–	✓
Multi-agent review	–	–	–	–	–	✓
End-to-end analysis	∼	✓	–	∼	✓	✓
Report generation	–	–	–	–	∼	✓
Minimal human oversight	✓	✓	✓	–	–	✓

Table 1: Comparison of agentic AI systems for HEP. ✓ = fully supported, ∼ = partially supported, – = not supported. "End-to-end analysis" refers to producing a complete measurement (selection through statistical inference); "minimal human oversight" indicates the system can run without continuous human feedback.

Productivity-enhancing tools or replacing physicists?

Caveats

Huge potential for HEP
Theory and Experiment, but

- AI agents may introduce subtle physics errors: strong need for human validation remains (is trade-off worth it?)
- AI agents less reliable for genuinely novel approaches, may ignores niche (but powerful) methods
- What is their complexity ceiling? Suitable for genuinely new analyses?
- Insufficiently robust review practices, danger of validation biases

Personal take: agentic AI is here to stay, and precisely for this fostering **critical and innovative thinking** is more important than ever

Summary and outlook

- Machine Learning and AI methods have revolutionised the way we do particle physics, both from the Theory and Experimental activities
- ML enables **major extensions of the physics program** of our colliders, such as Higgs self-coupling discovery at the LHC
- Agentic AI can streamline many tasks through **end-to-end pipelines**, but does not replace (actually, makes it more important than ever) **critical human thinking**
- As HEP defines the post-LHC era, ML/AI will be **ubiquitous at any future collider**