Accelerated AI Algorithms for Data-Driven Discovery


# Teaching an old experiment new tricks

**Phil Harris**

Phil Harris Lab@MIT

# AI as a positive disrupter

- We are all at a meeting to discuss the prospects of AI
  - It is remarkable what we can do with it
    - We can now prototype software, do analysis very quickly
  - We should be excited about the prospects
- This is an opportunity to teach an old experiment new tricks
    - Time to remove/revitalize the old frozen software frameworks
    - Can we change the sociology revolving around AI
      - We don’t need less people, we need more physics

# Just one Example

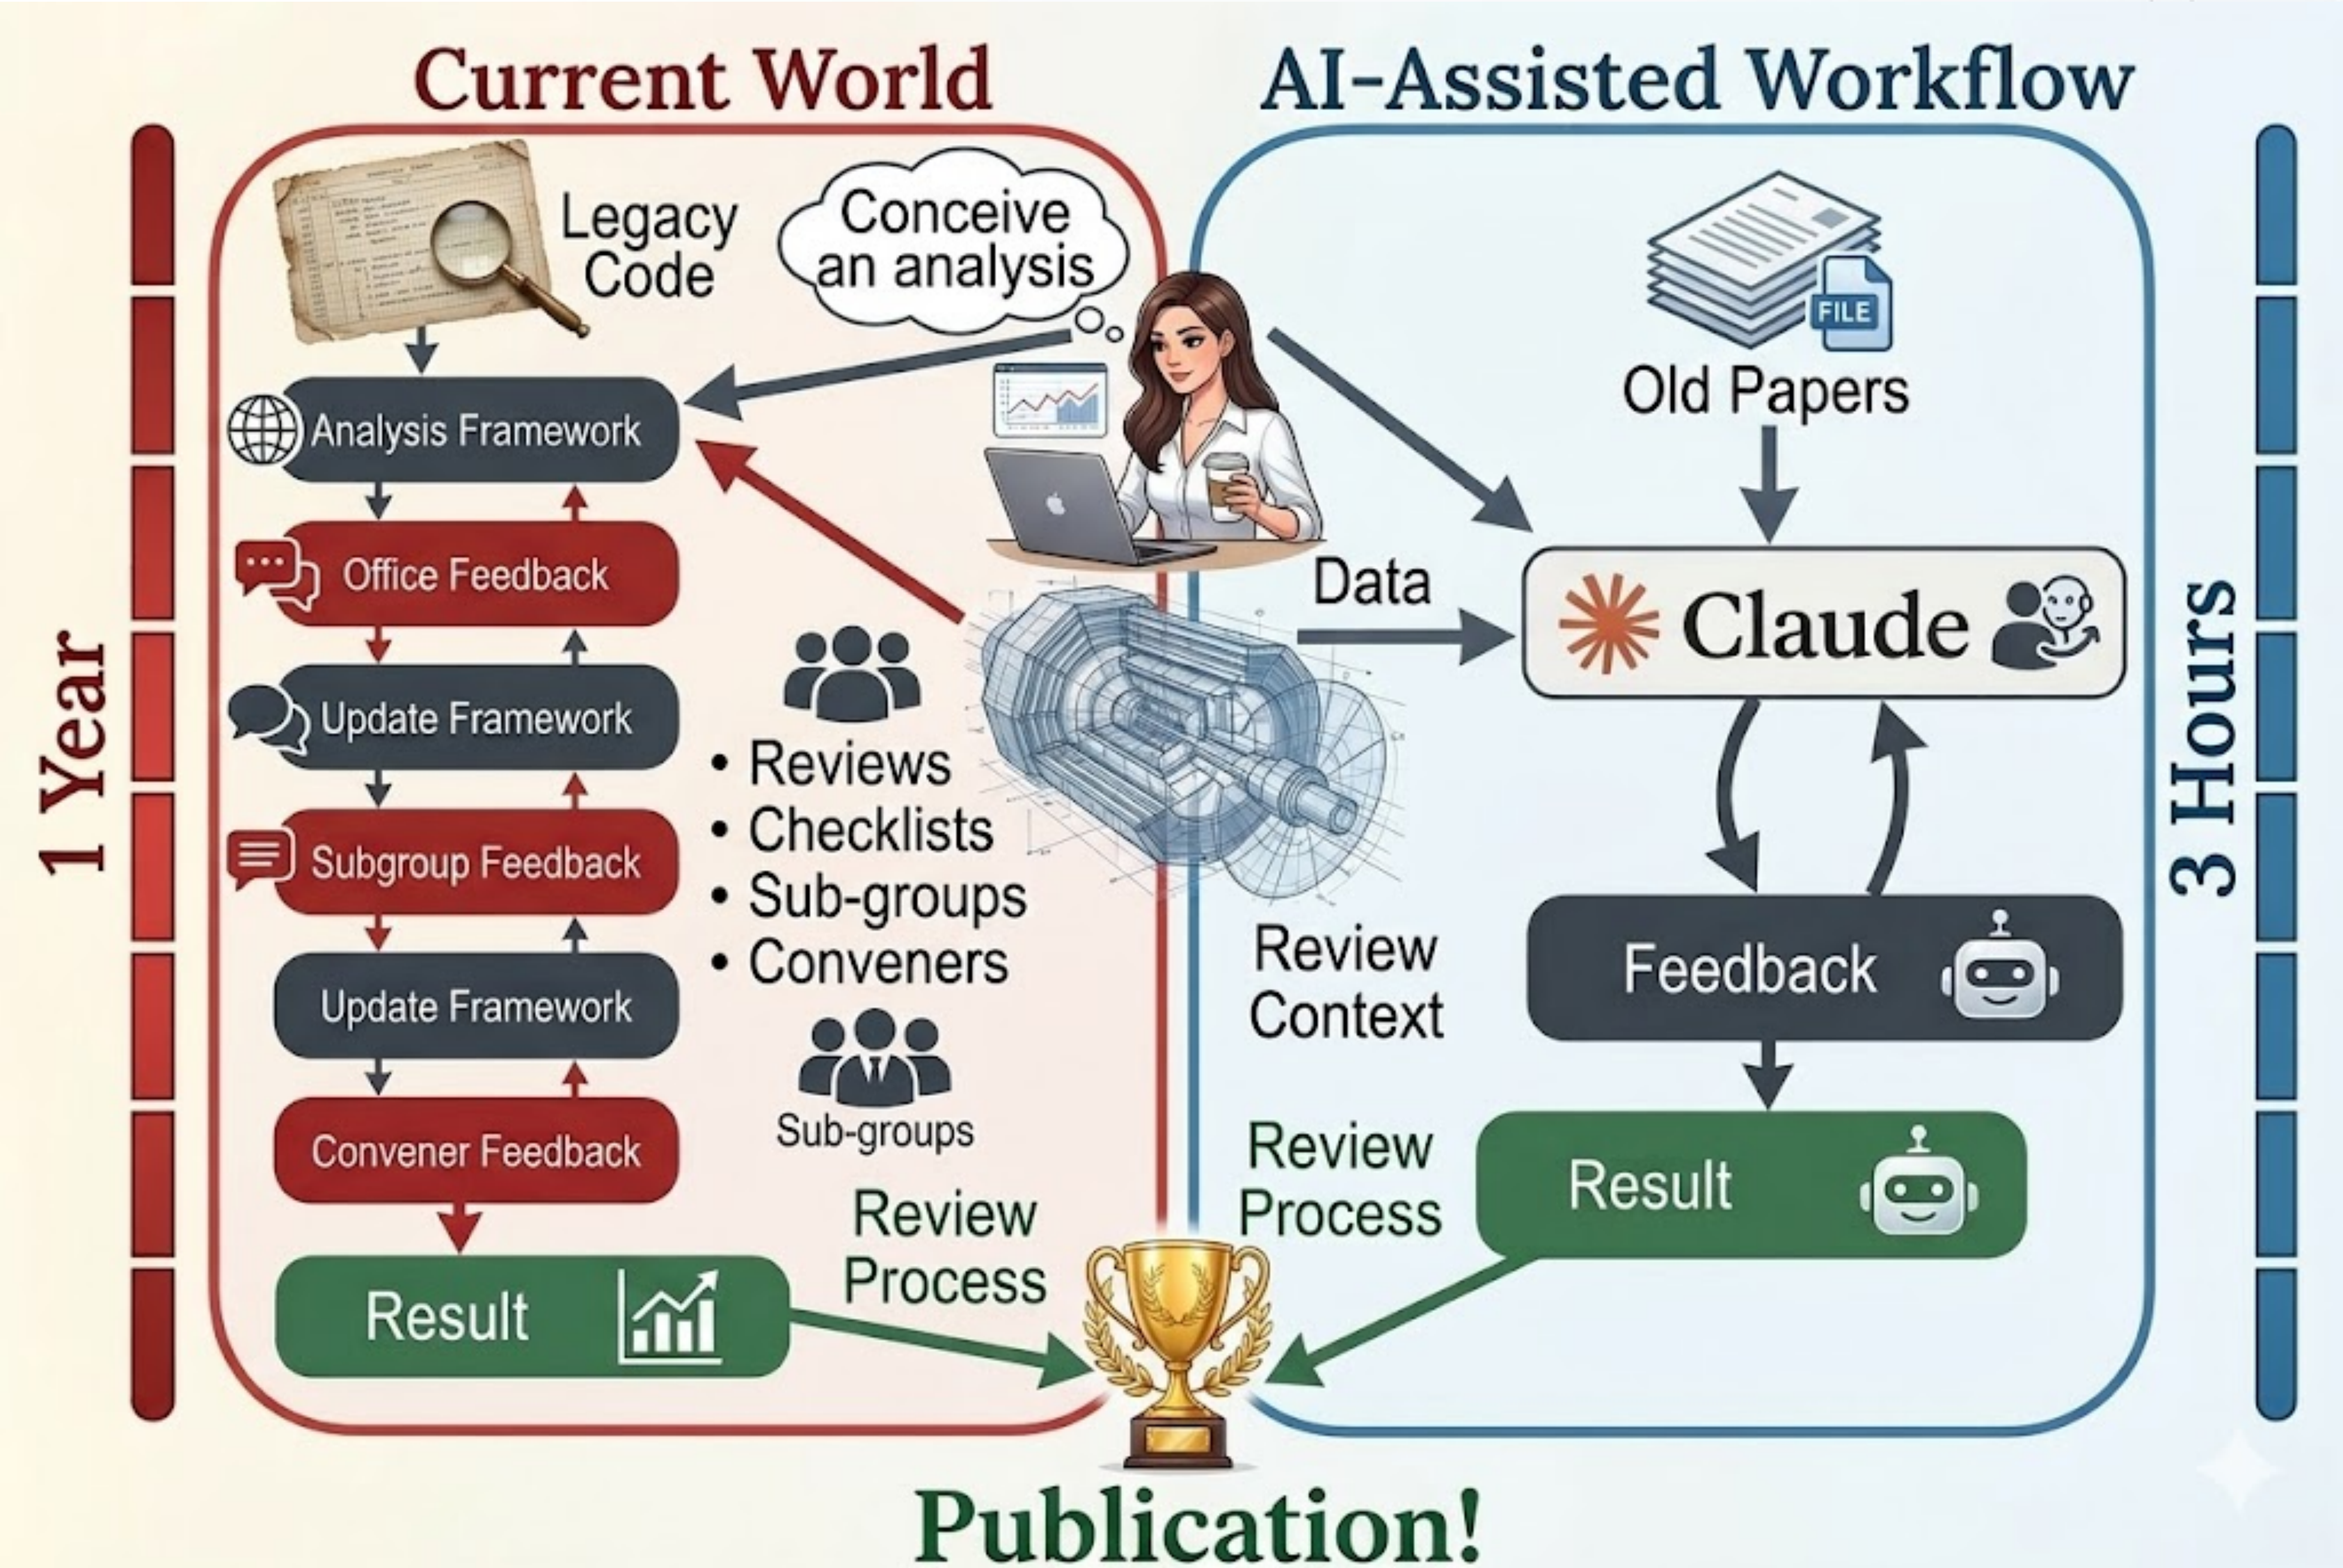

# What does this mean?

- We can think bigger and bolder about our research
  - Less analysis frameworks/Streamlined review process
  - Preservation of Physics data becomes easy
- Ultimately students/PD/Profs go back to physics
  - Students will work more like professors

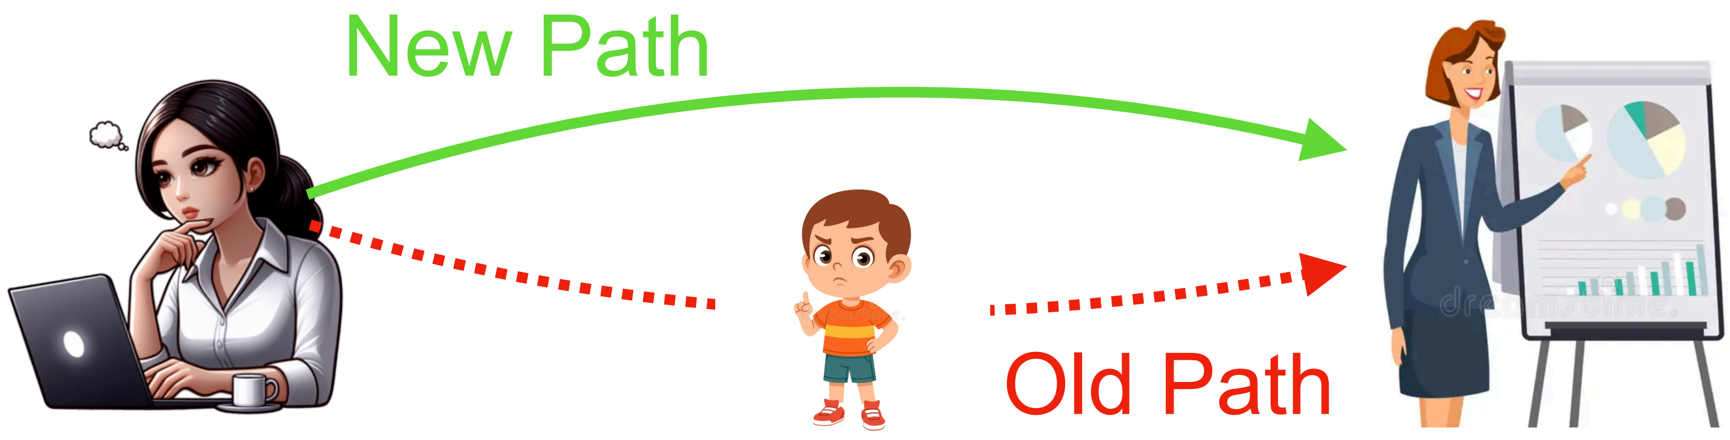

# Current Successes


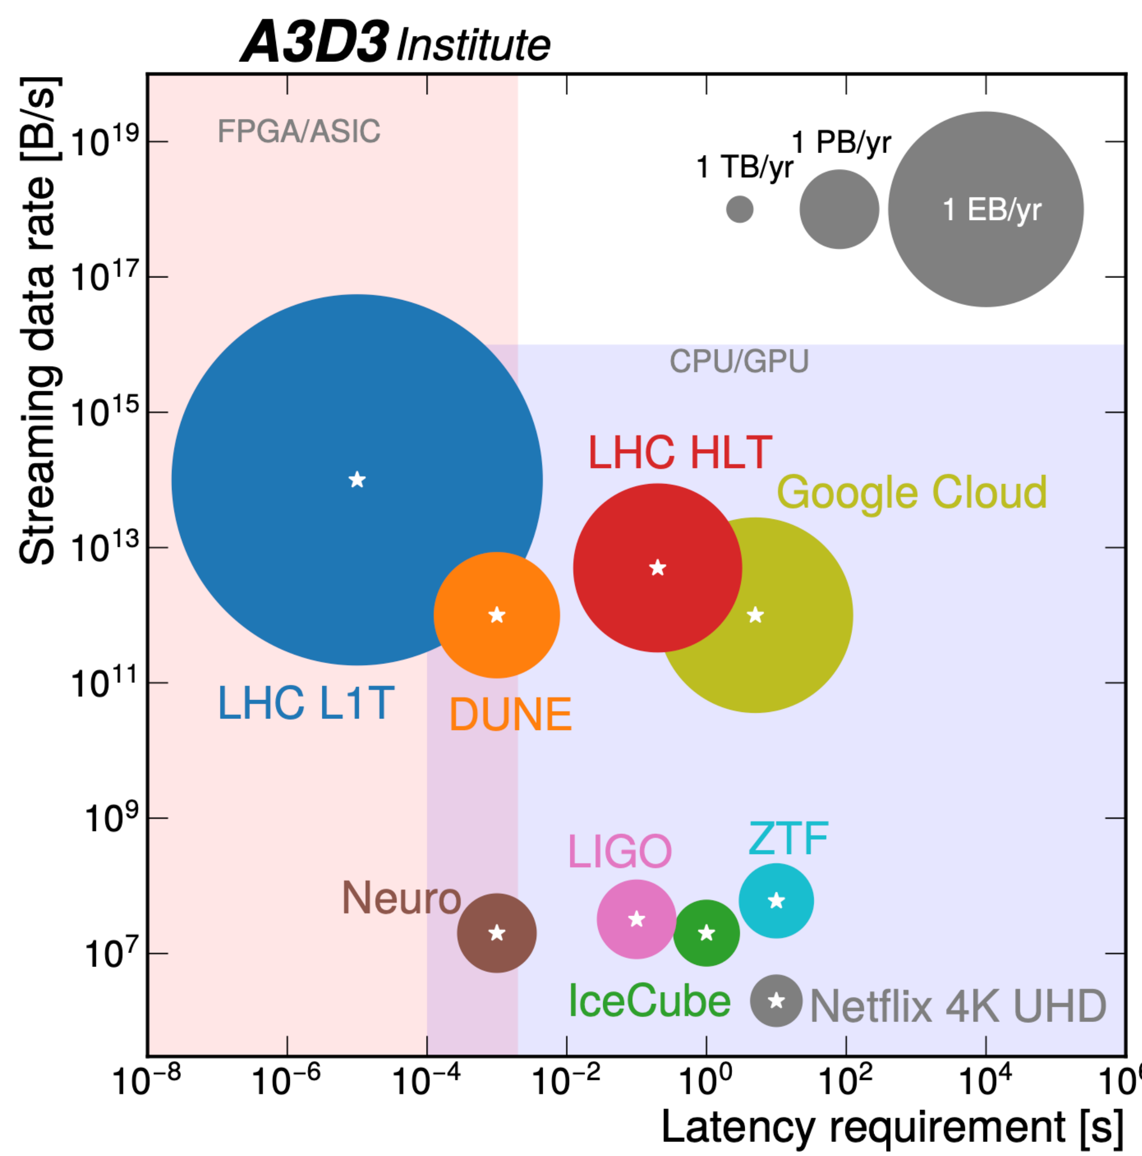


We have managed to present a new problem that engaged industry and othe rcollaborations

# A few ideas

- We are not alone here :
  - LIGO/Dune/CMS/ATLAS/Astro/??? Need same tools
  - Large computing facilities are ubiqutous
  - There is a large community excited by high energy physics
    - Interesting problems with open data can engage everybody
- Could we rewrite from scratch Athena/CMSSW?
  - Can we make sure all production is fully automated?
  - Can we make it completely GPU native?

# How do we do this

New AI initiative for the LHC

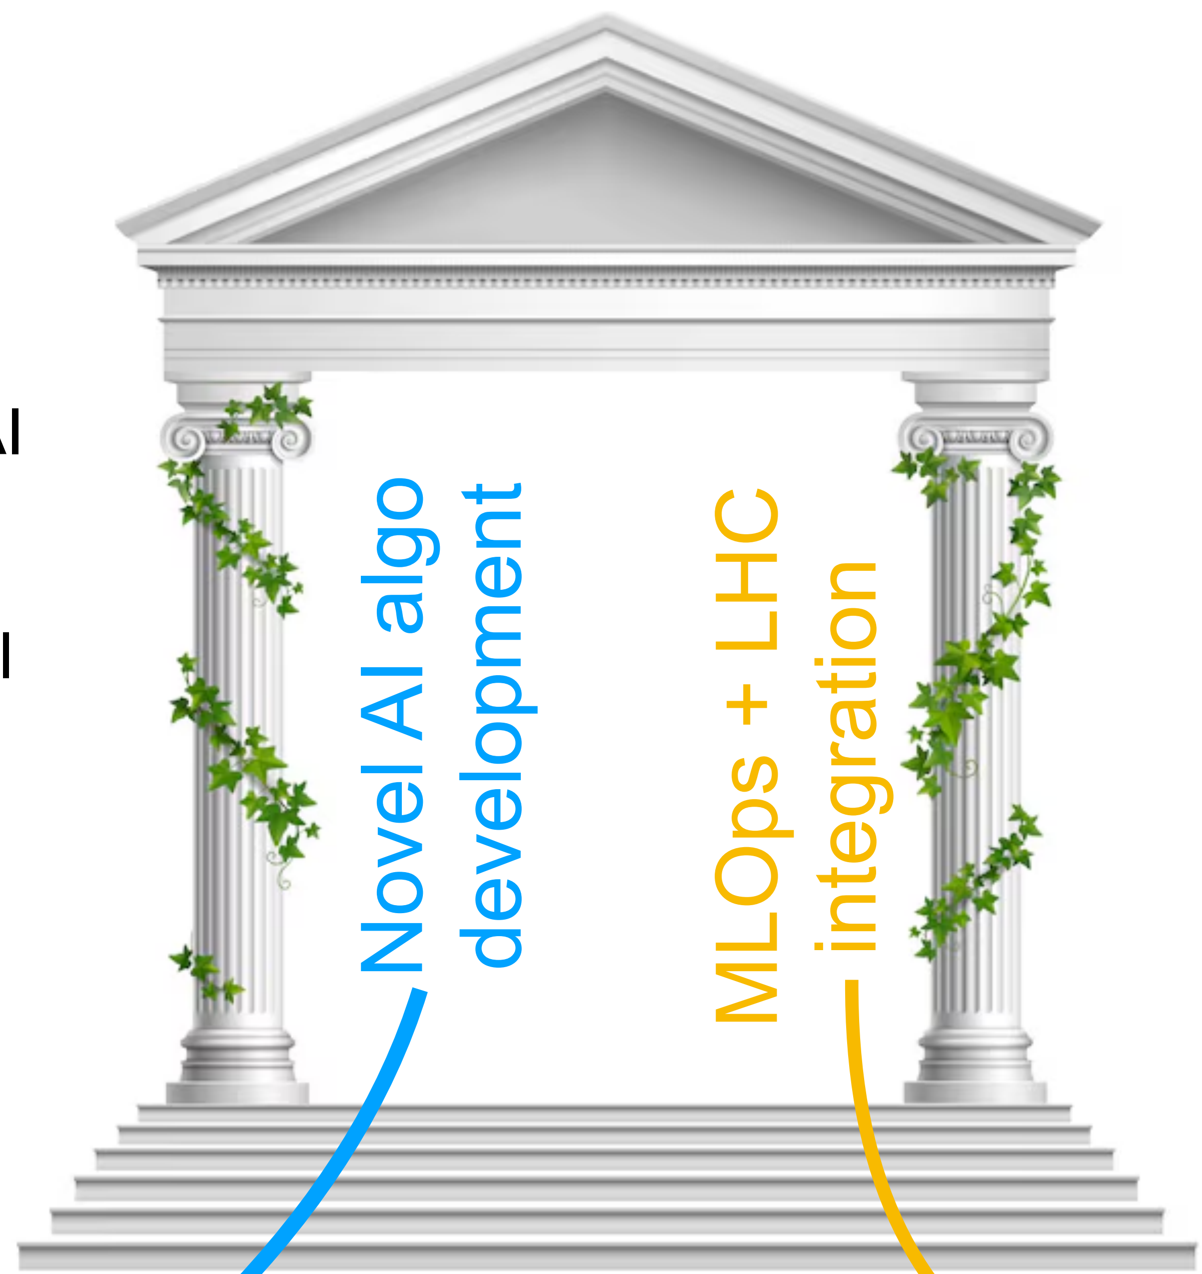


Development of AI benchmarks

Novel multi-modal scaling laws

Validating AI for science

Deployment of updated AI software

New agentic analysis

Automated workflows

Novel approaches and visibility within the AI community

New Physics Capabiities


# Thanks!

# Pursuing full automation of Physics Analysis

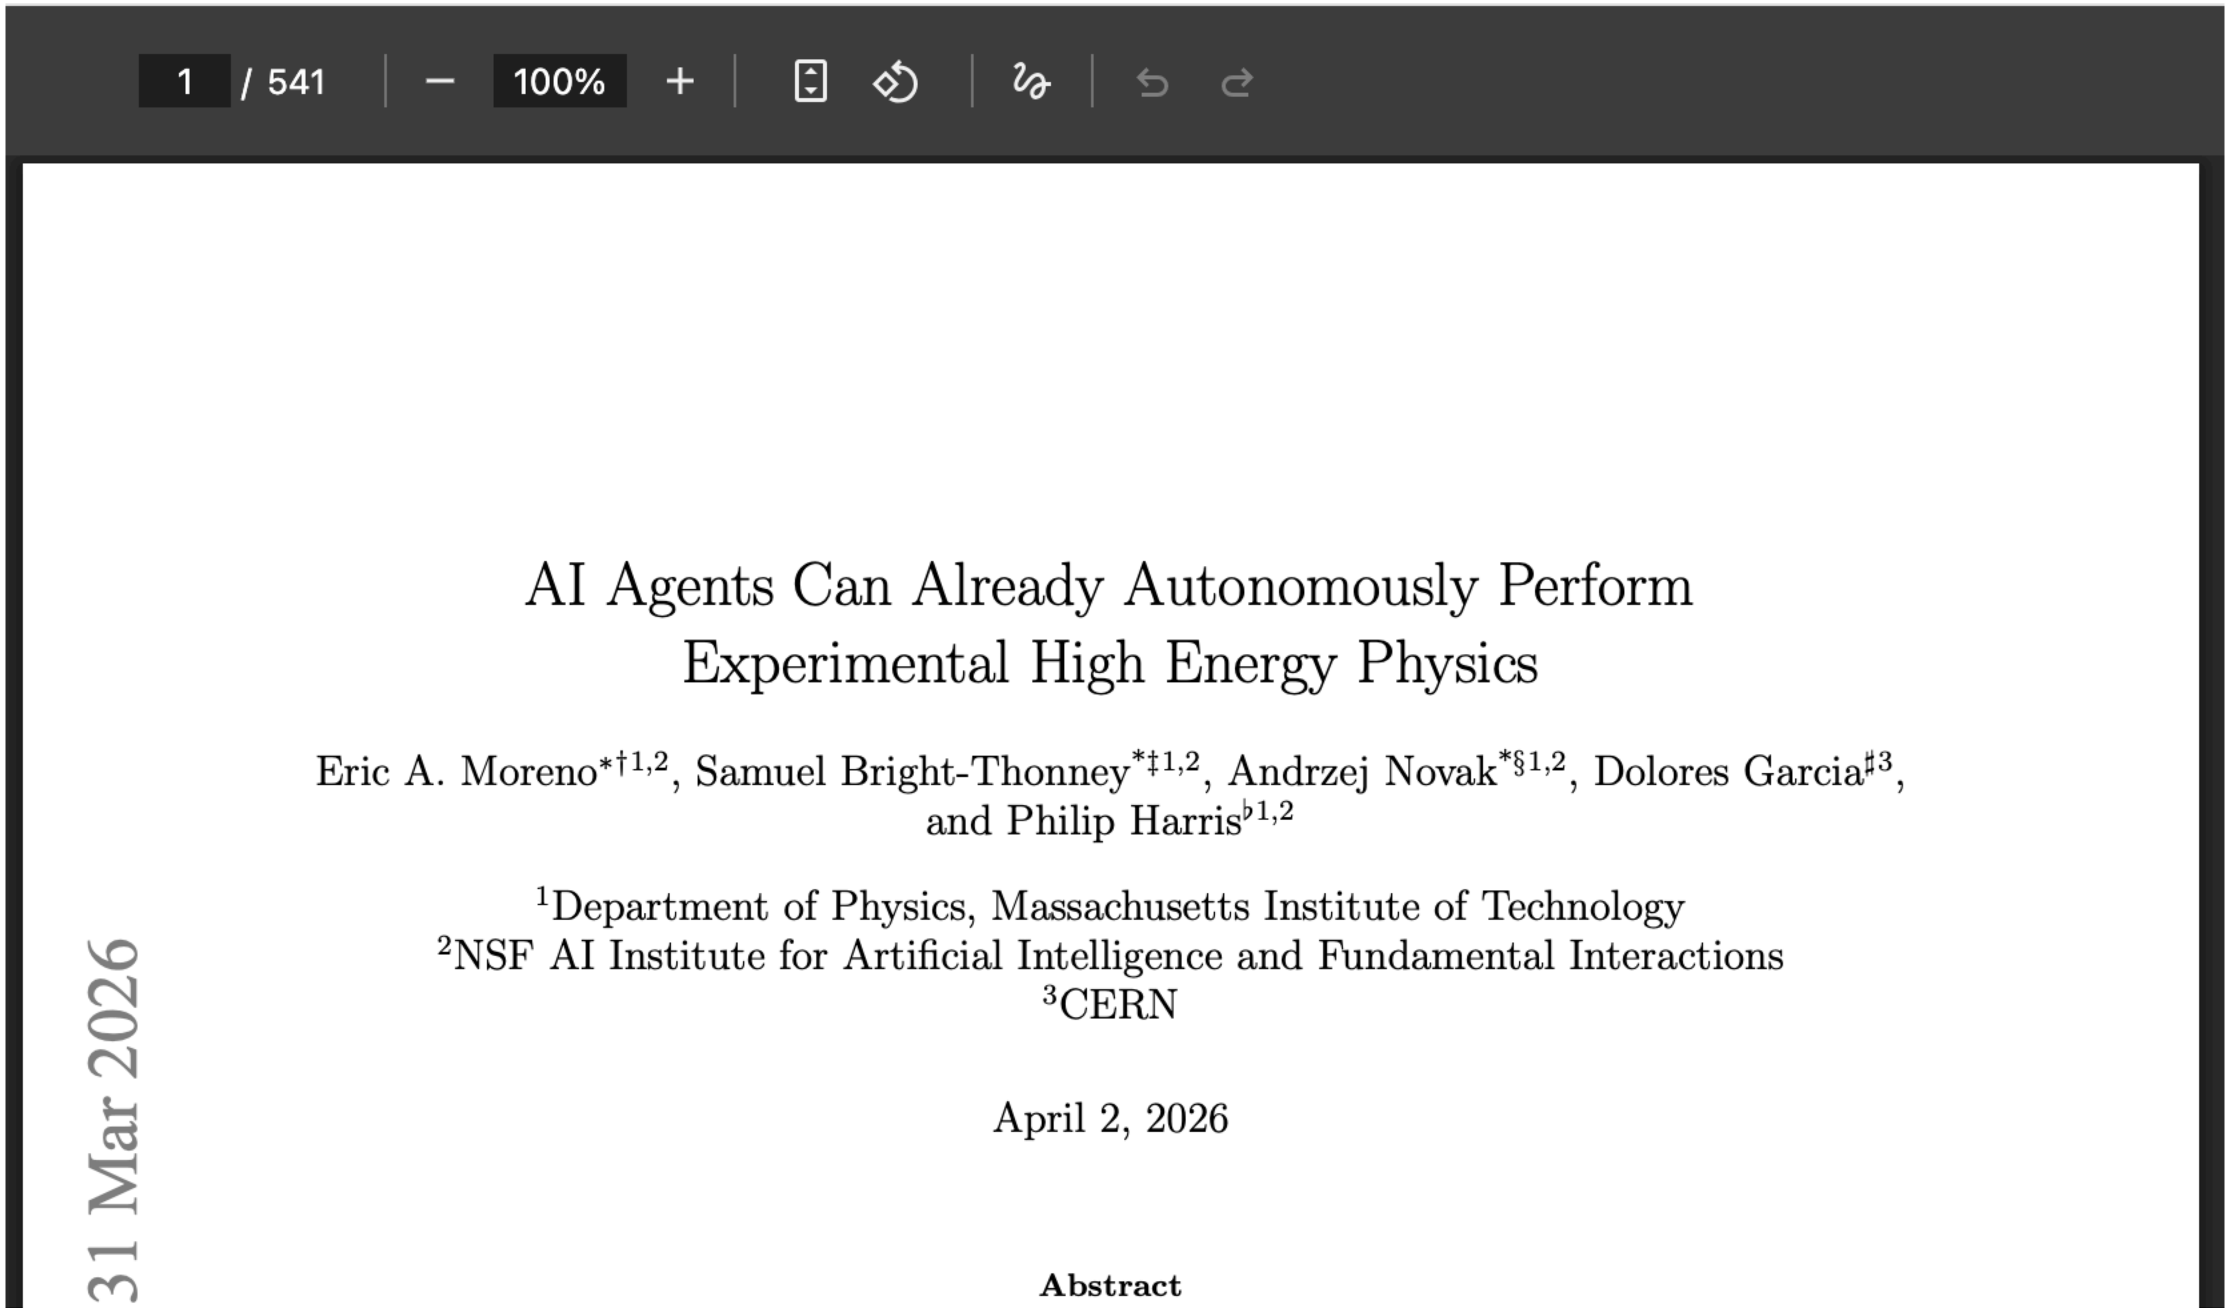
1 / 541 100%

31 Mar 2026

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]

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

April 2, 2026

**Abstract**

https://arxiv.org/pdf/2603.20179

# How good is it?

- This is just from two days
- Its at 4-th year grad student level

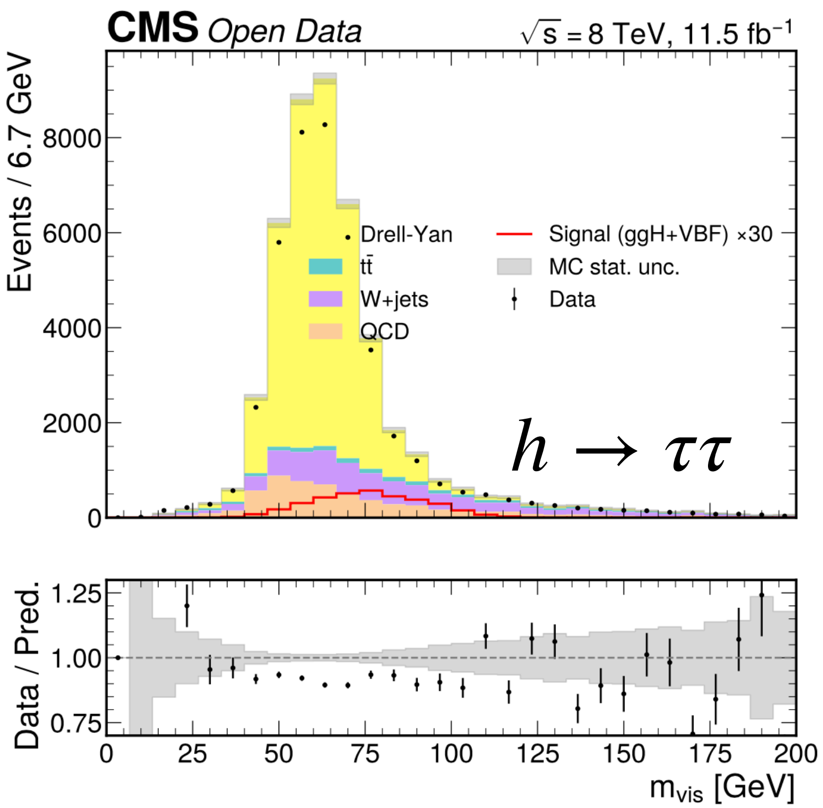


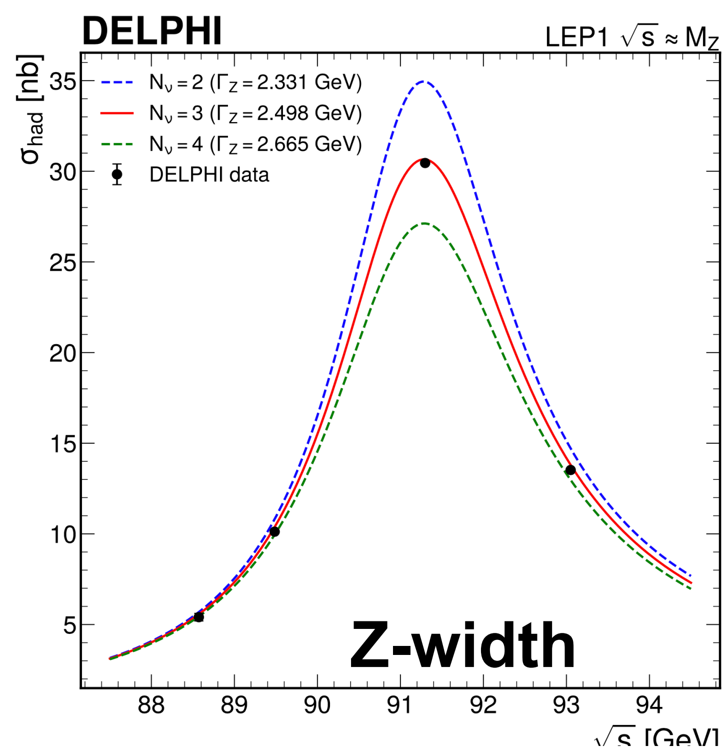


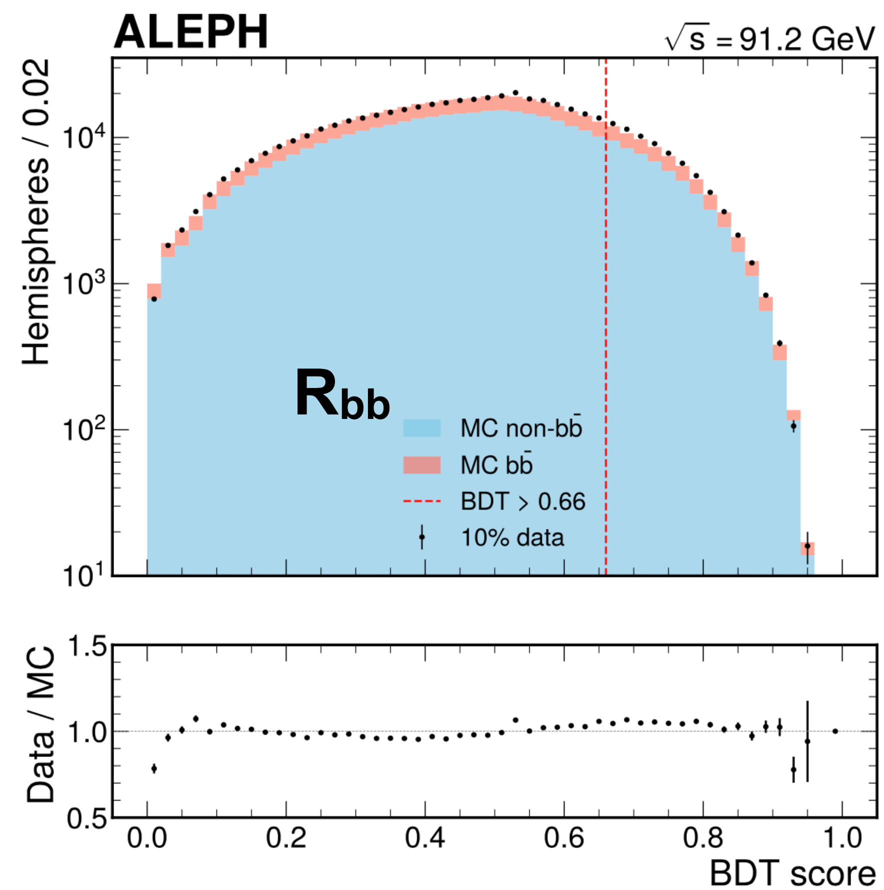


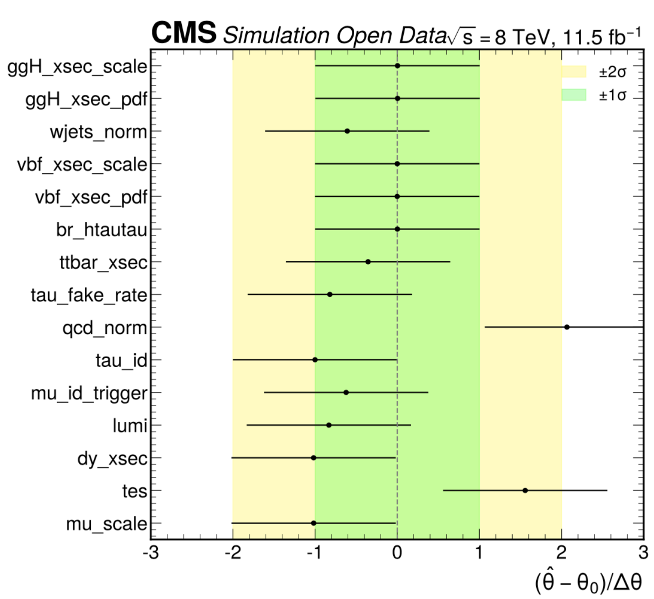


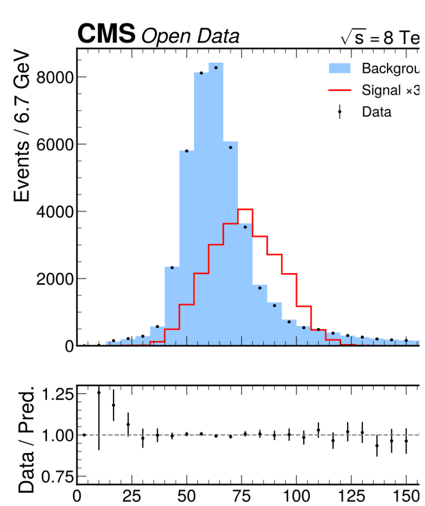


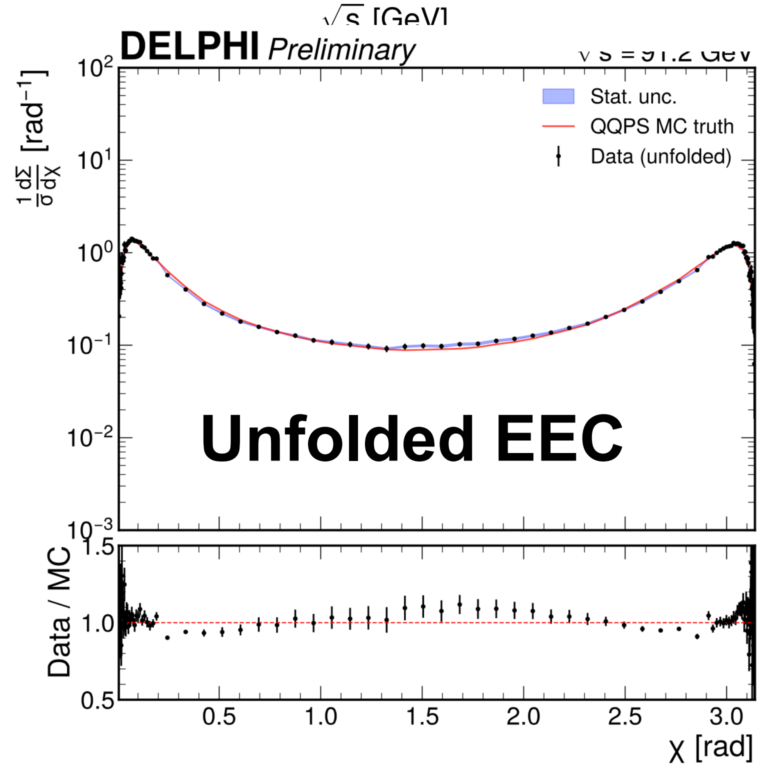


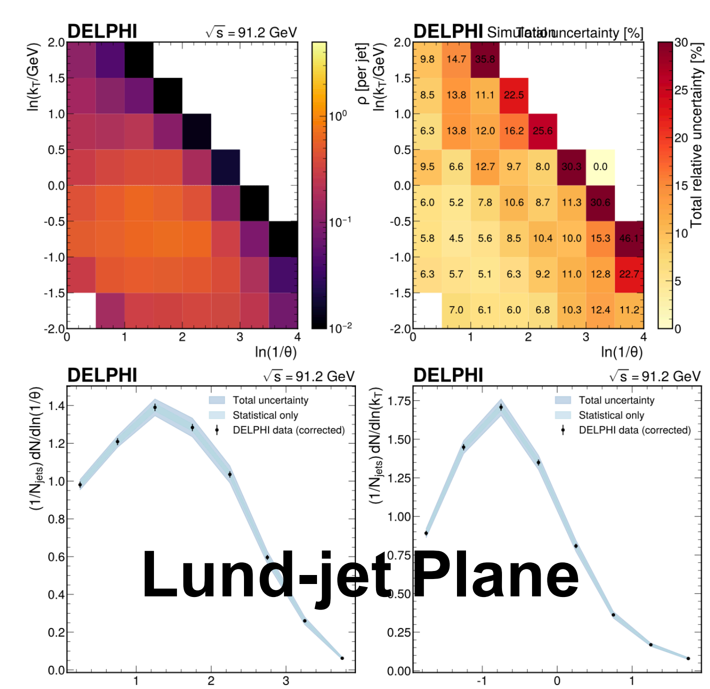

LHC datasets
Have become the
standard for low power
AI benchmarks

# Current Successes

| Dataset | Model | Accuracy (%) | LUT | FF | DSP | BRAM | $F_{max}$ (MHz) | Latency (ns) | Area×Delay (LUT×ns) |
|---|---|---|---|---|---|---|---|---|---|
| JSC CERNBox | **KANELÉ** | **75.1** | **5034** | 1917 | 0 | 0 | **870** | 8.1 | $\mathbf{4.1 \times 10^4}$ |
| | NeuraLUT-Assemble [6] | 75.0 | 8539 | 1332 | 0 | 0 | 352 | **5.7** | $4.87 \times 10^4$ |
| | AmigoLUT-NeuraLUT [47] | 74.4 | 42742 | 4717 | 0 | 0 | 520 | 9.6 | $4.10 \times 10^5$ |
| | PolyLUT-Add [30] | 75.0 | 36484 | 1209 | 0 | 0 | 315 | 16 | $5.84 \times 10^5$ |
| | NeuraLUT [5] | **75.1** | 92357 | 4885 | 0 | 0 | 368 | 14 | $1.29 \times 10^6$ |
| | PolyLUT [4] | 75.0 | 246071 | 12384 | 0 | 0 | 203 | 25 | $6.15 \times 10^6$ |
| | LogicNets [42] | 72.0 | 37931 | **810** | 0 | 0 | 427 | 13 | $4.93 \times 10^5$ |
| JSC OpenML | **KANELÉ** | 76.0 | **1232** | 900 | 0 | 0 | **987** | 7.1 | $8.7 \times 10^3$ |
| | NeuraLUT-Assemble [6] | 76.0 | 1780 | 540 | 0 | 0 | 941 | **2.1** | $\mathbf{3.92 \times 10^3}$ |
| | TreeLUT [23] | 75.6 | 2234 | **347** | 0 | 0 | 735 | 2.7 | $6.03 \times 10^3$ |
| | DWN [7] | 76.3 | 4972 | 3305 | 0 | 0 | 827 | 7.3 | $3.6 \times 10^4$ |
| | da4ml [37] | **76.9** | 12250 | 1502 | 0 | 0 | 212 | 18.9 | $2.3 \times 10^5$ |
| | hls4ml (Fahim et al.) [16] | 76.2 | 63251 | 4394 | 38 | 0 | 200 | 45 | $2.85 \times 10^6$ |
| MNIST | **KANELÉ** | 96.3 | 3809 | 4133 | 0 | 0 | 864 | 9.3 | $3.5 \times 10^4$ |
| | NeuraLUT-Assemble [6] | **97.9** | 5070 | 725 | 0 | 0 | 863 | **2.1** | $\mathbf{1.06 \times 10^4}$ |
| | TreeLUT [23] | 96.6 | 4478 | **597** | 0 | 0 | 791 | 2.5 | $1.12 \times 10^4$ |
| | DWN [7] | 97.8 | **2092** | 1757 | 0 | 0 | 873 | 9.2 | $1.92 \times 10^4$ |
| | PolyLUT-Add [30] | 96.0 | 14810 | 2609 | 0 | 0 | 625 | 10 | $1.48 \times 10^5$ |
| | AmigoLUT-NeuraLUT [47] | 95.5 | 16081 | 13292 | 0 | 0 | **925** | 7.6 | $1.22 \times 10^5$ |
| | NeuraLUT [5] | 96.0 | 54798 | 3757 | 0 | 0 | 431 | 12 | $6.58 \times 10^5$ |
| | PolyLUT [4] | 97.5 | 75131 | 4668 | 0 | 0 | 353 | 17 | $1.38 \times 10^6$ |
| | FINN [43] | 96.0 | 91131 | — | 0 | 5 | 200 | 310 | $2.82 \times 10^7$ |
| | hls4ml (Ngadiuba et al.) [33] | 95.0 | 260092 | 165513 | 0 | 345 | 200 | 190 | $4.94 \times 10^7$ |

We have managed to present a new problem that engaged industry and othe rcollaborations

# What does a Prompt Look Like?

Scaffold and run a measurement analysis of the primary Lund jet plane density in hadronic Z decays using archived ALEPH data at $\sqrt{s} \sim 91.2$ GeV.

Setup: scaffold analyses/lund_plane as a measurement, set data_dir=[...] in .analysis_config, install the pixi environment, then begin orchestrating.

Observable: The 2D density of primary Cambridge/Aachen declusterings in each thrust hemisphere, mapped to coordinates $(\ln 1/\Delta\theta,\ \ln k_t/\mathrm{GeV})$, where $\Delta\theta$ is the emission angle and $k_t = E_{\mathrm{soft}}\, sin\Delta\theta$. Use charged particles only (**`pwflag`** $==$ 0, **`highPurity`** $==$ 1). One jet = one hemisphere.

Deliverables:

1. 2D density $\rho(\ln 1/\Delta\theta, \ln k_t)$ corrected to charged-particle level, 10–15 × 10–15 bins at least - as fine as it makes sense
2. 1D projections ($k_t$ spectrum, angular spectrum) with covariance
3. Number of primary declusterings vs. minimum $k_t$ threshold
4. Comparison to PYTHIA 6 MC
5. Machine-readable results (CSV/NPY)

Data: [...]

Think of each agent as human(student) that sp physics who you give instructions to in a bunc text files

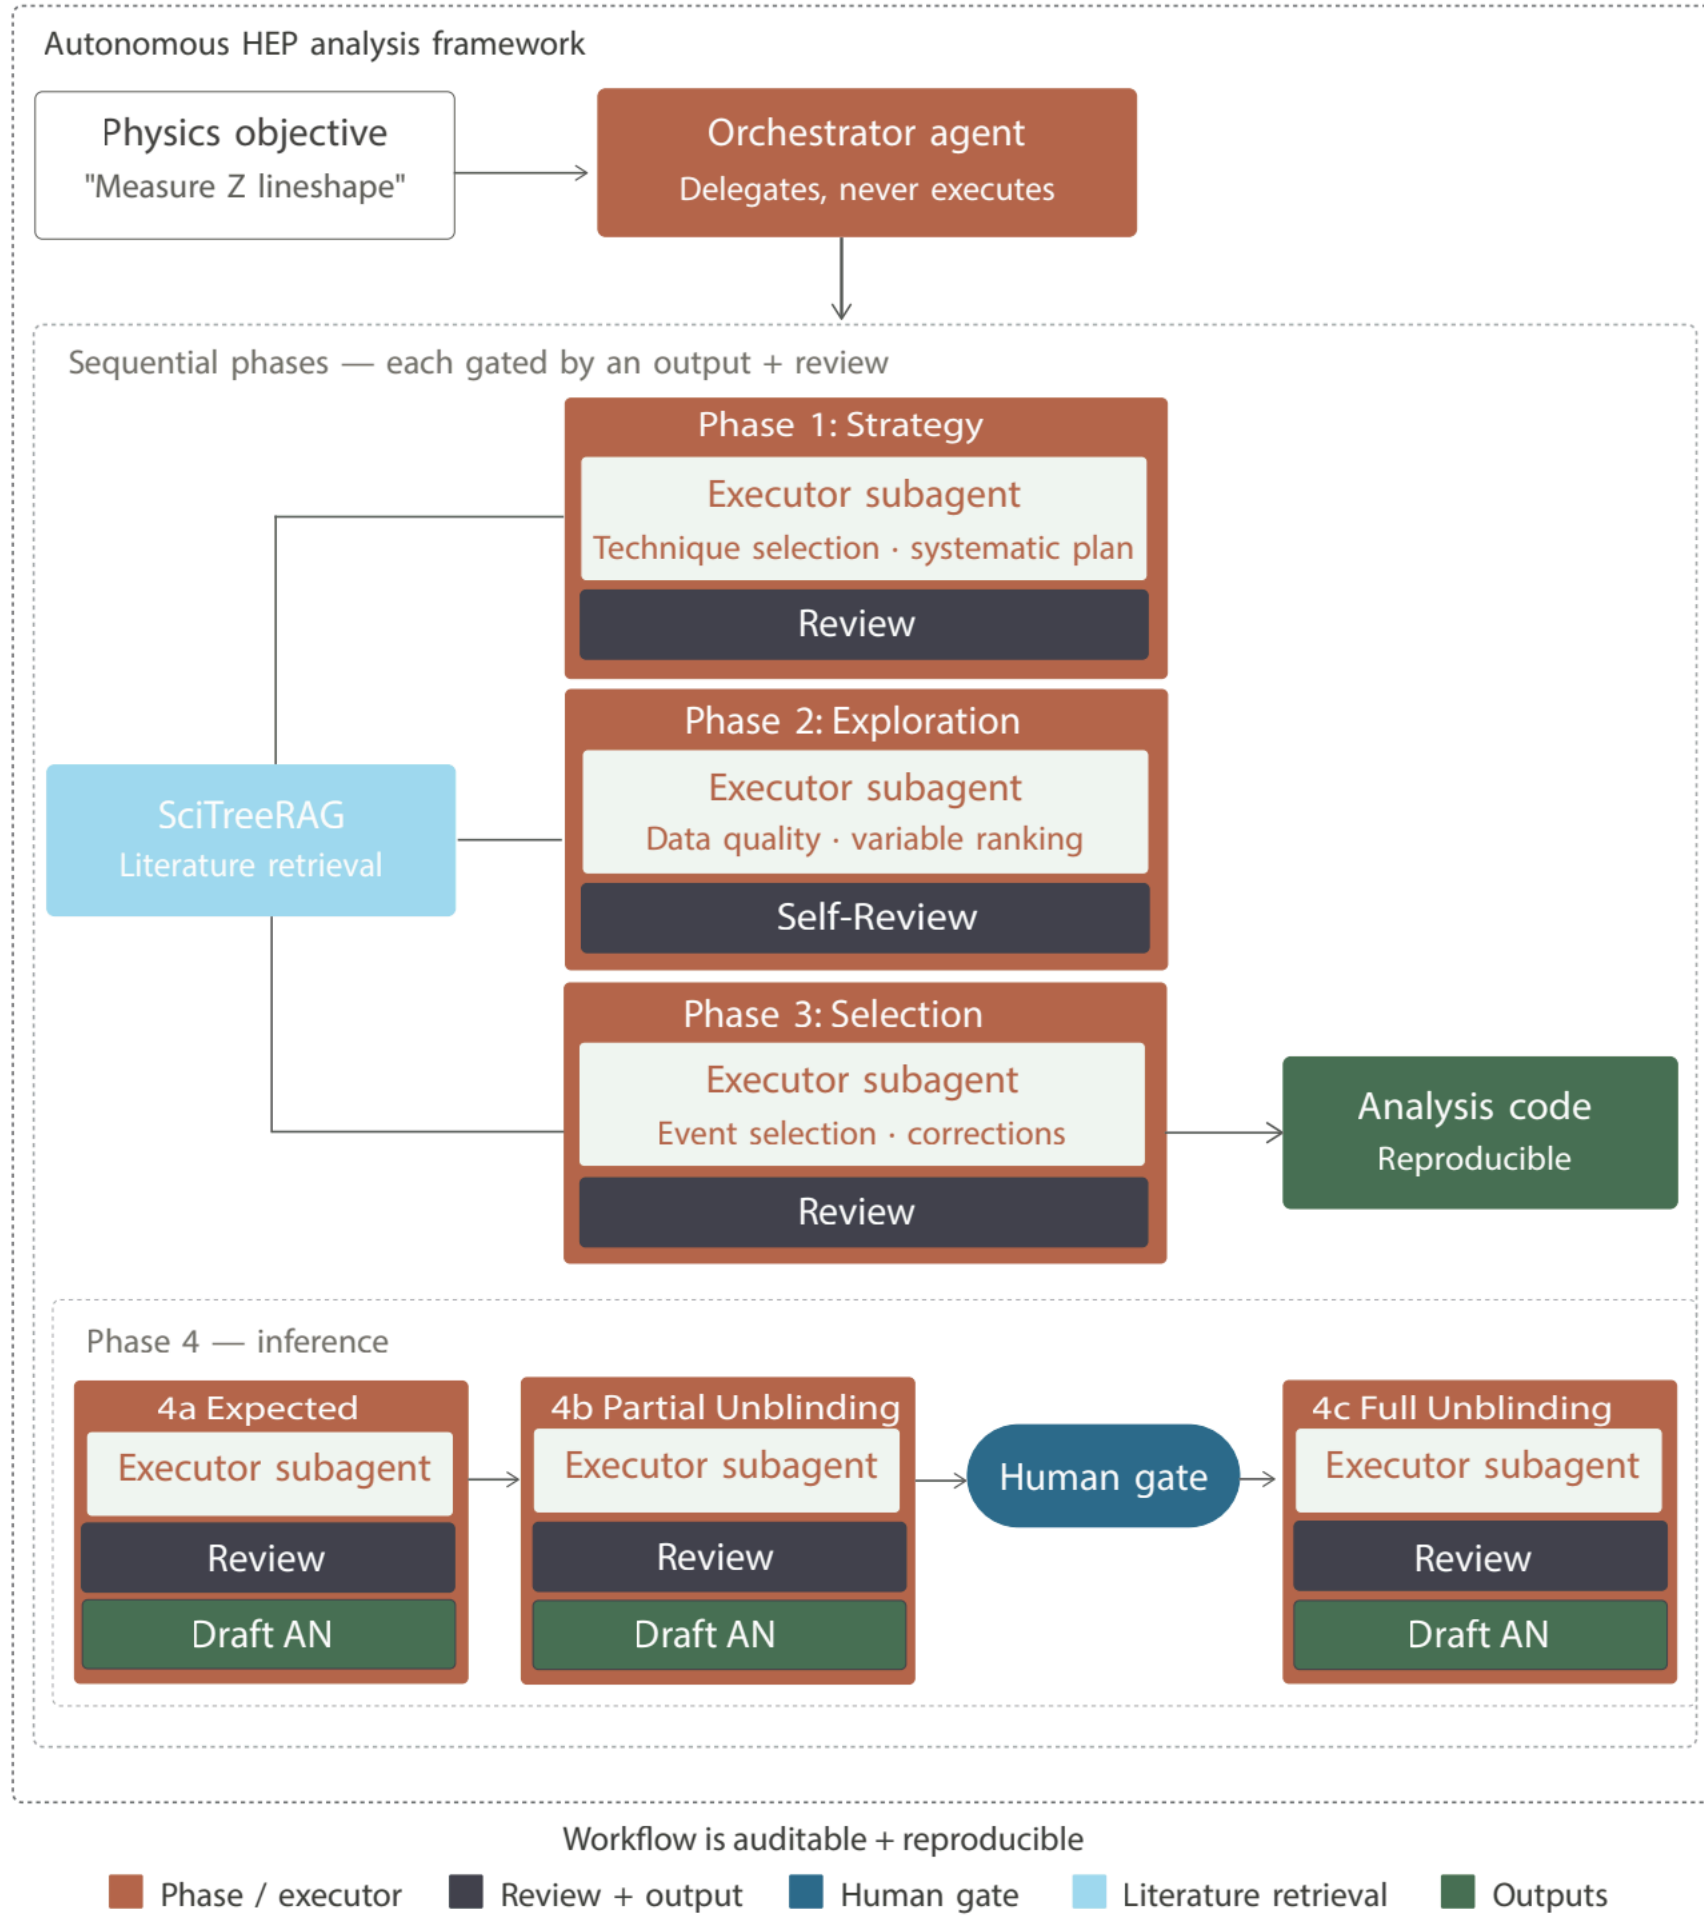

# How good is it really?

## How does it compare to a reference

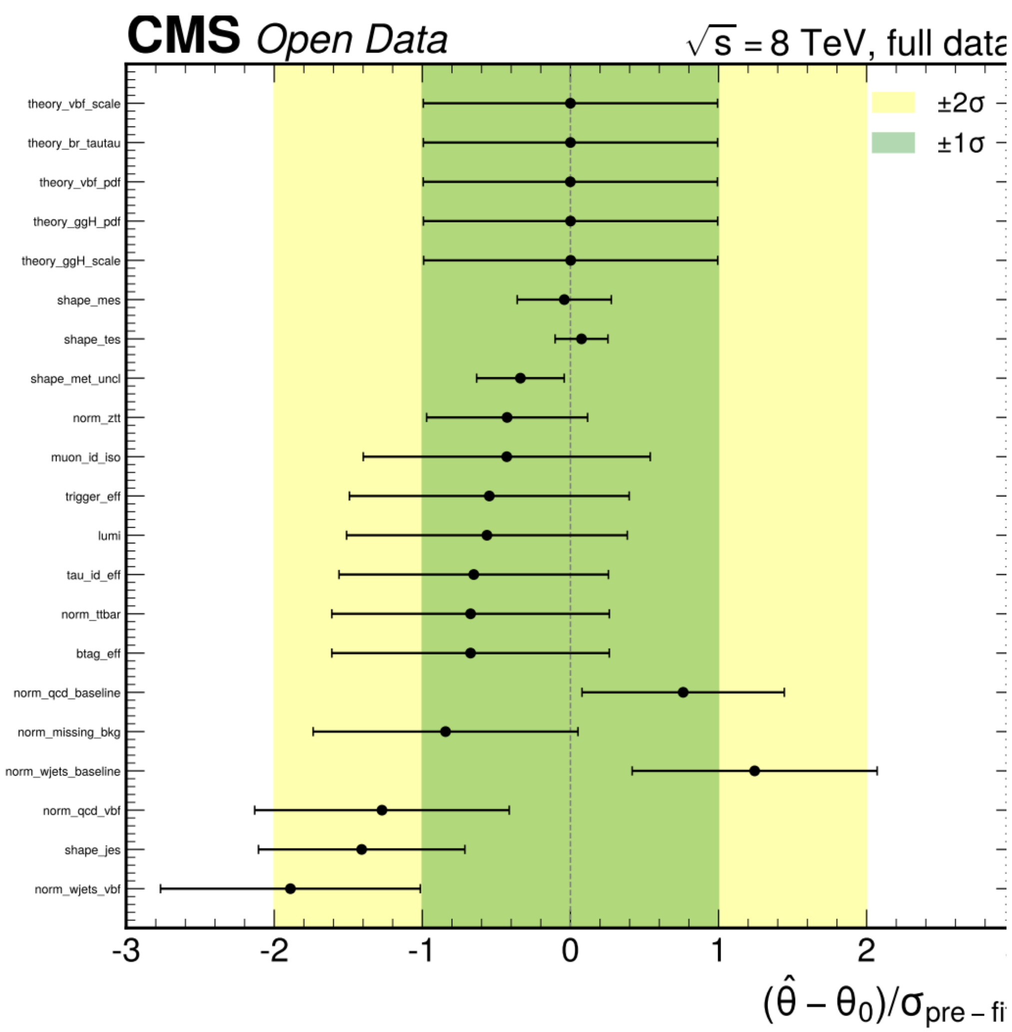


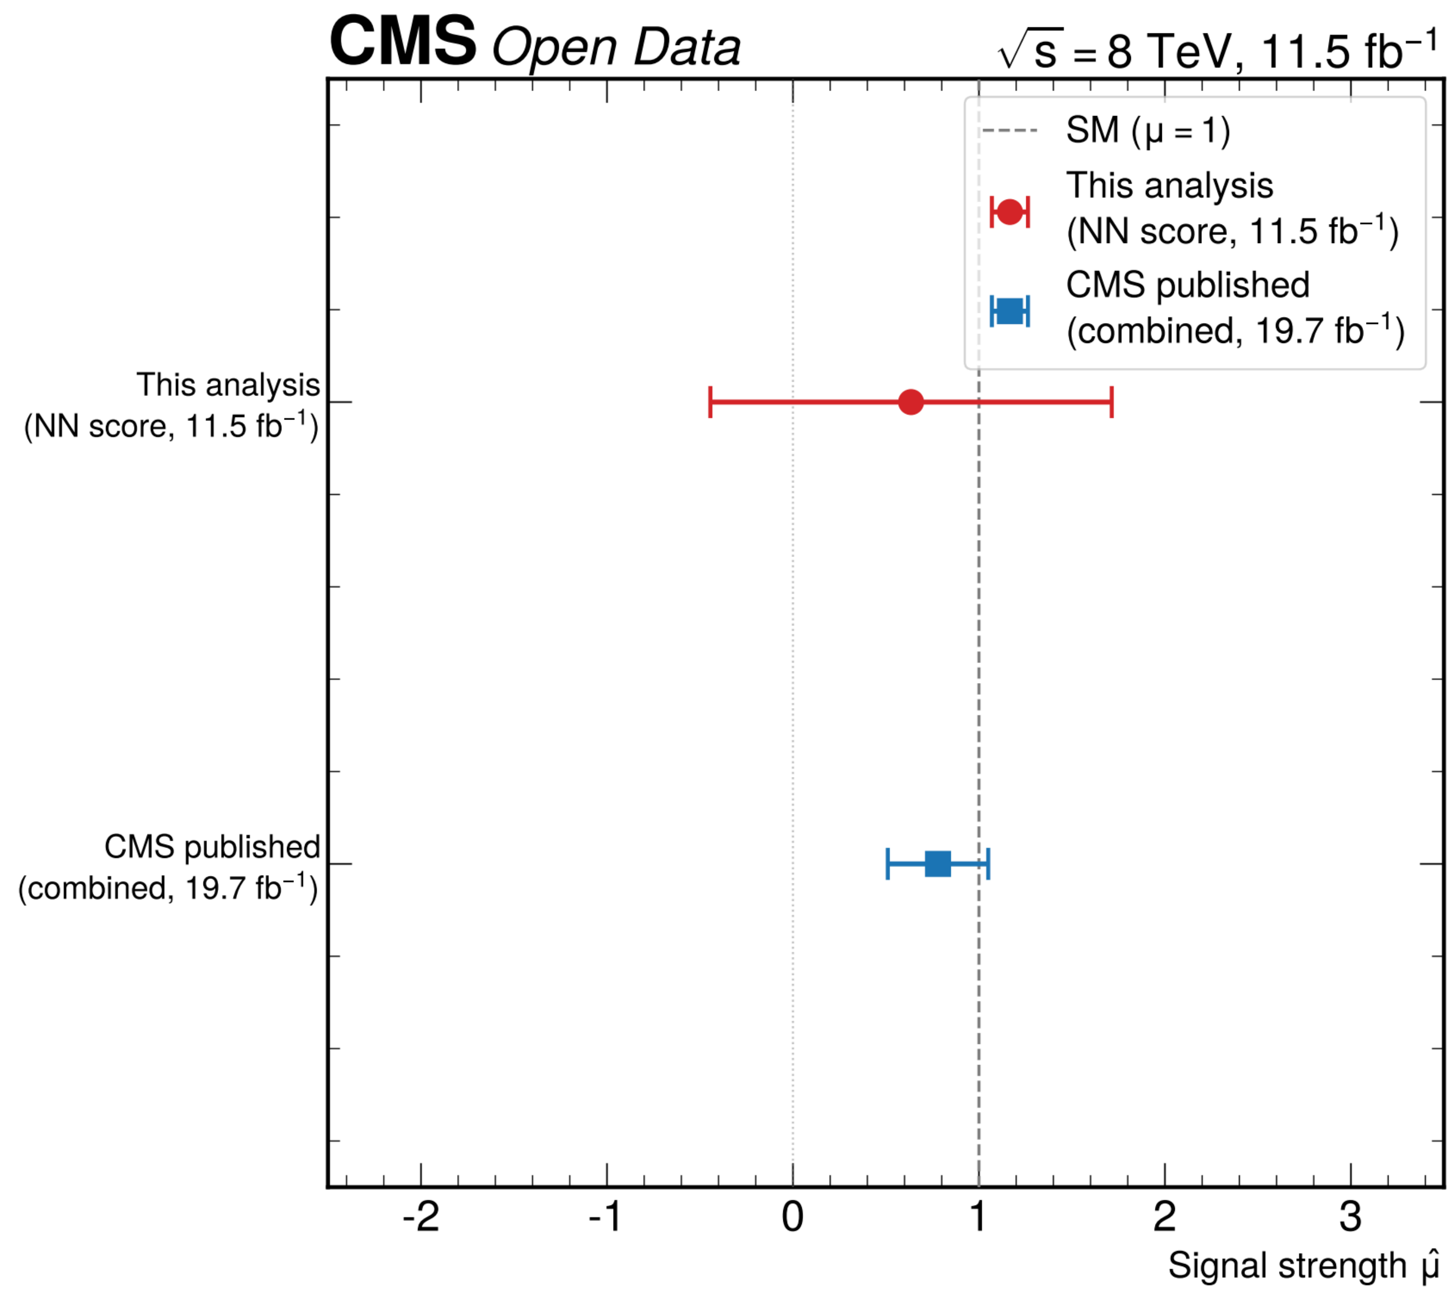


LLM: 0.63 ±1.08

CMS*: 1.01±0.59

ATLAS*:1.1±1.09

Same Uncertainty as the ATLAS experiment

*Very approximate scaling

# Come and Join FastML !

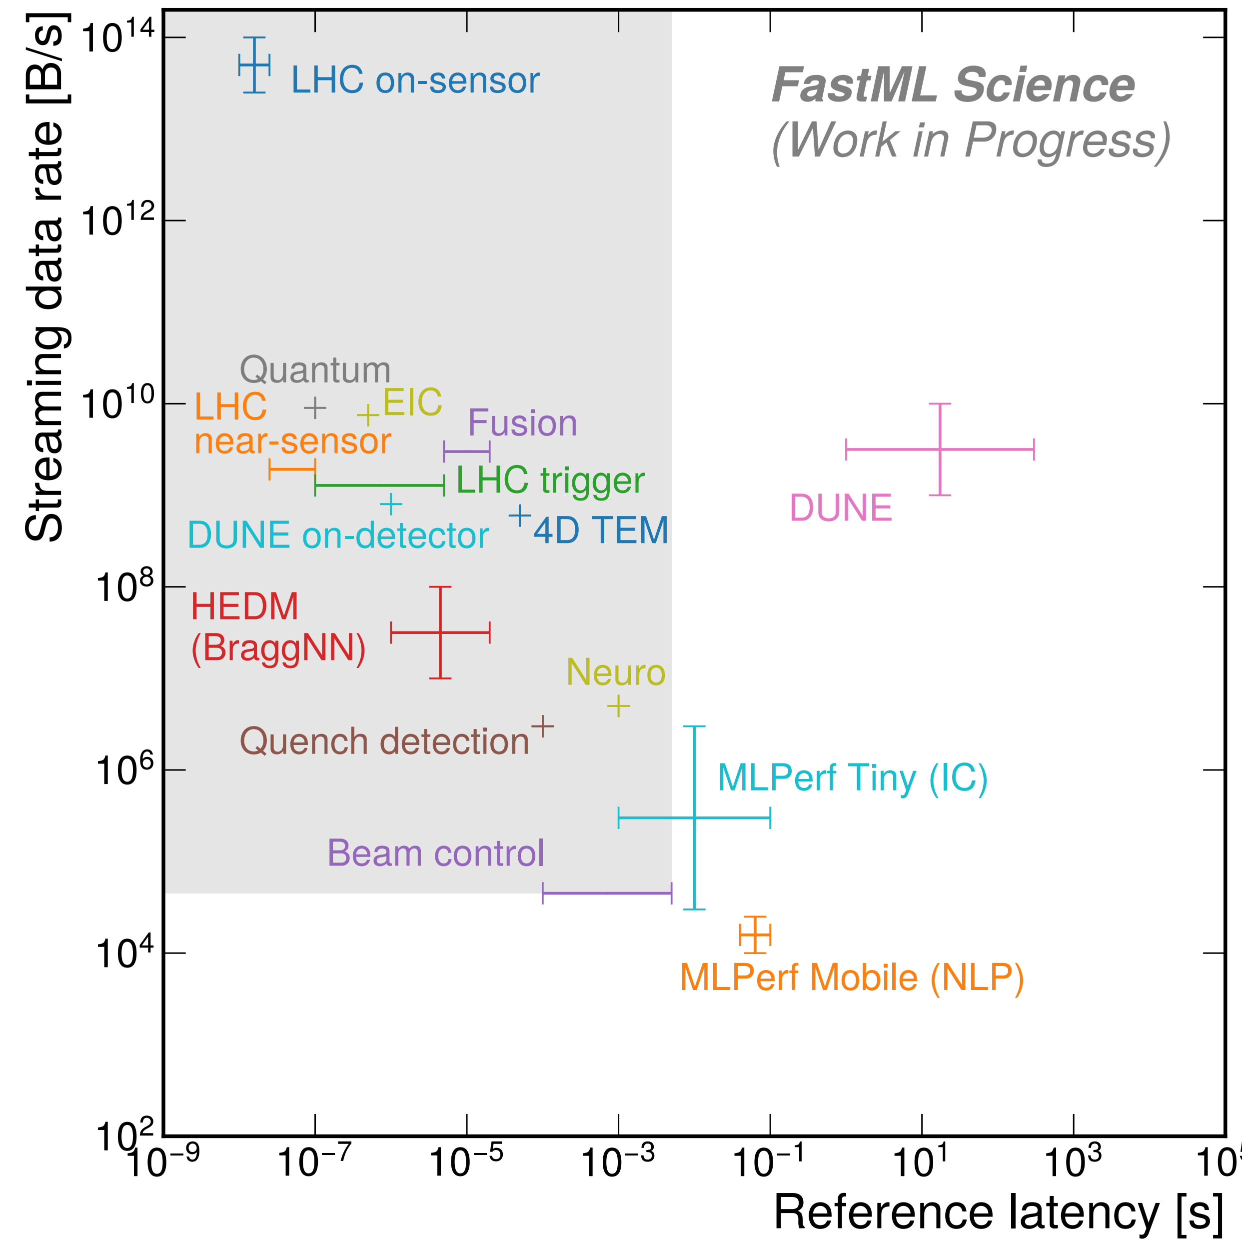


Opening doors for real-time AI in HEP and many other fields