UNIVERSITY OF LIVERPOOL

AI & (Q)Computing & Software work and R&D

Eduardo Rodrigues, for many colleagues (see names in pages)

Particle Physics Annual Meeting, Liverpool, 21st – 22nd May 2026
https://indico.ph.liv.ac.uk/event/2385/

Quantum Computing

See their talks …

Samuel Godwood, Marina Maneyro

- Paper **"Fault-Tolerant Resource Comparison of Qudit and Qubit Encodings for Diagonal Quadratic Operators"** https://arxiv.org/abs/2604.26792

- Developing and comparing qudit and qubit algorithms for quantum-simulation for lattice field theories
- Recent work focuses on fault-tolerant resource estimates, identifying regimes where qudits may offer benefits savings over qubit-based implementations

- Early fault-tolerant algorithm applications
- Using sampling quantum circuits to understand physical systems of interest, with a focus on nuclear-model spectra and ground-state energies

Fault-Tolerant Resource Comparison of Qudit and Qubit Encodings for Diagonal Quadratic Operators

Samuel Godwood,[1,*] Doğa Murat Kürkçüoğlu,[2] Gabriel N. Perdue,[2] Marina Maneyro,[1] and Alessandro Roggero[3,4]

[1] *Department of Physics, University of Liverpool, Liverpool, L69 7ZL, United Kingdom*
[2] *Fermi National Accelerator Laboratory, Batavia, Illinois 60510, USA*
[3] *Dipartimento di Fisica, University of Trento, Via Sommarive 14, I-38123 Povo, Trento, Italy*
[4] *INFN-TIFPA Trento Institute of Fundamental Physics and Applications, Via Sommarive 14, I-38123 Povo, Trento, Italy*

Finite local Hilbert-space truncations arise naturally in quantum simulations of lattice field theories and motivate qudit encodings, but their fault-tolerant advantage over qubit encodings remains unclear. We compare the non-Clifford cost of implementing quadratic diagonal evolutions, exemplified by $U = e^{-it\phi_x^2}$ in a uniform field-amplitude discretization of a real scalar field, using either one logical d-level qudit or $n_b = \lceil \log_2 d \rceil$ logical qubits. We analyze two standard settings: product-formula simulation and LCU/block encoding, taking the resource metric to be the number of non-Clifford gates after synthesis into a discrete logical gate set. Because tight synthesis bounds for general single-qudit rotations are not known, we express the qudit constructions in terms of embedded two-level $SU(2)$ rotations and derive explicit finite-d break-even conditions for their synthesis cost; these serve as compiler targets for when qudit encodings can outperform the qubit baseline. Within the constructive models studied here, product-formula implementations would require an exponentially stronger per-primitive synthesis advantage for qudits to win asymptotically, while in the LCU setting the qubit encoding is asymptotically cheaper in d. Nevertheless, the finite-d threshold analysis identifies low dimensional regions in which qudits can yield meaningful constant-factor savings, particularly for LCU-based implementations. As a secondary analysis of the LCU construction, we use an idealized negligible-overhead qubit-qudit code-switching model to give an absolute T-count comparison, and reinterpret the savings as an allowable per-switch overhead budget.

Sampling quantum circuits to reconstruct nuclear-model energy spectra.

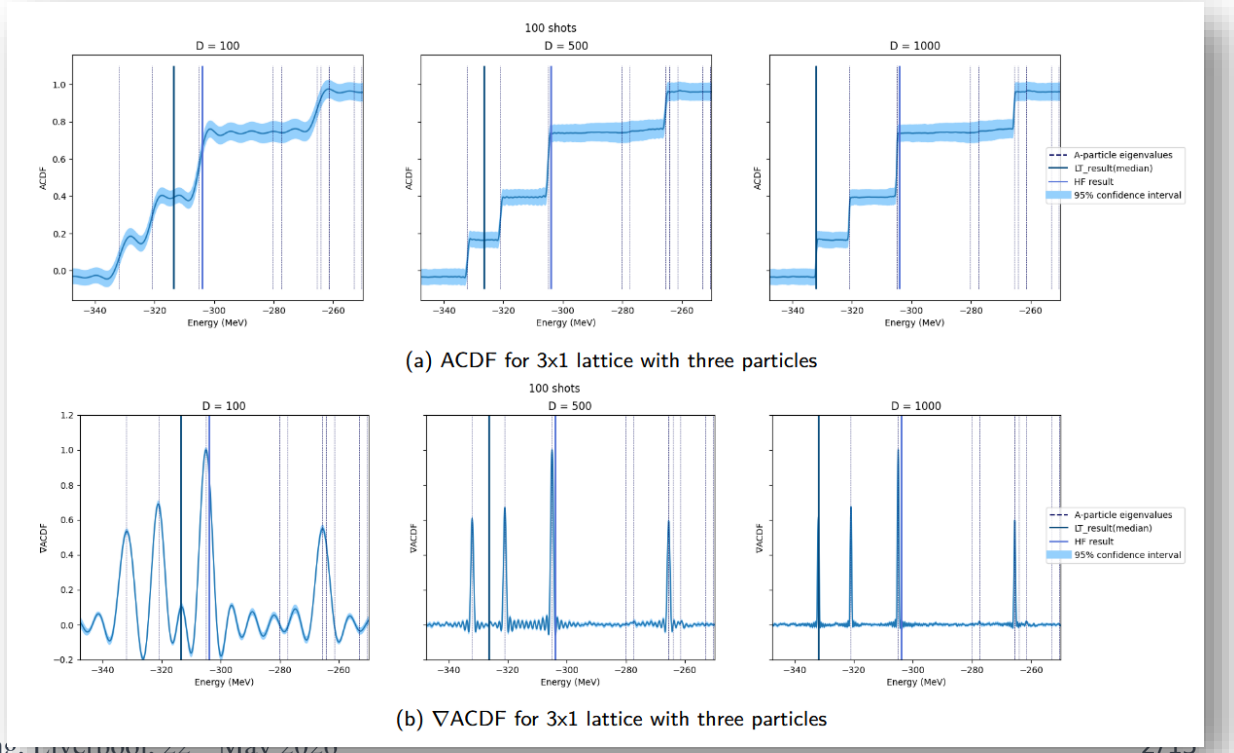

ATLAS – AI/ML in physics analyses

Several analyses using ML (tools) – most relevant contributions from Liverpool are:

- **Physics object**: b-jet, c-jet tagging (Nikos, Andy)
  and **tau-hadronic reconstruction** (Jordy, Monica, Mehul, and Lennox + Rob in previous years)
- **Di-Higgs analysis in bbtautau** (Carl, Jordy, Buphesh, Lennox)
- **Higgs invisible** (Stephen, Monica, Andy)
- **Higgs into long-lived axion-like particles** (Oliver, John, Monica, Nikos)
- **Compressed-scenario SUSY searches** (John, Monica + had MPHYS student this year working on this, and there is a new PhD coming in October 2026 shared with RAL)

ATLAS presentation from yesterday talks about several of these analyses …

LHCb – LLMs for FPGA firmware development

- Exploring the use of LLMs for VHDL/Verilog
- Using both paid and free models
  - Even with paid models - **open weights preferred**
  - 50/50 running locally vs. cloud
  - Stealing time on **RTX 5060 Ti 16GB** @ CERN
- Very rapid pace of change
  - In 9 months, models went from “jank” to “usable”
  - A lot to learn
- FPGA development is still very niche
  - Nowhere near the generated code of Python or JS
- **What’s it good for?**
  - Testbenches and simulations
  - Boilerplate
  - Documentation/Explanation
- **What’s it bad at?**
  - Timing (laughably bad - “just add 1 clock cycle delay idiot”)

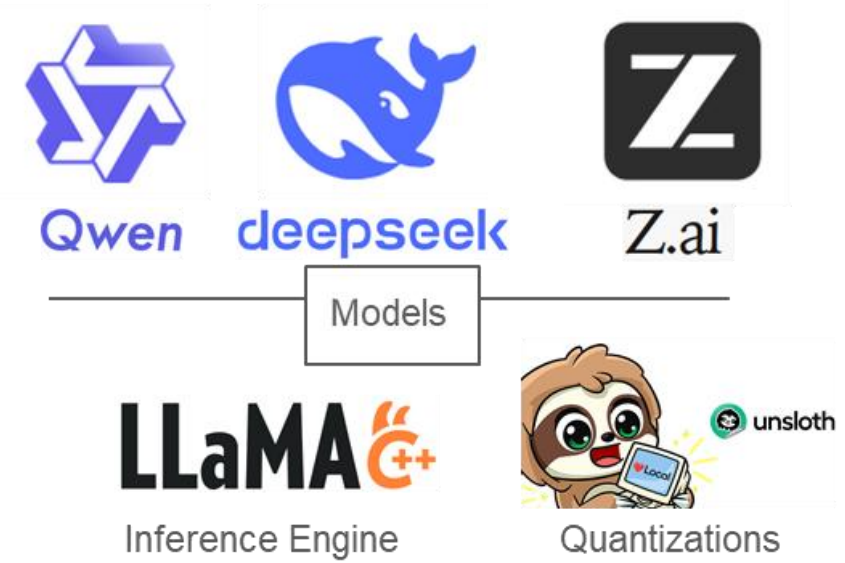

Pi coding harness

LHCb – ML for tracking

Giuseppe Costantino, Karol, Kurt, David

- **ML for tracking is a big, important and interest topic**
- Liverpool LHCb investigating with the VELO as a use case – Giuseppe's PhD

See Giuseppe's talk from yesterday …

- **Prototype algorithm for real-time reconstruction of particle tracks in the VELO detector runs on GPUs**
  (exploit CUDA-based parallelism)
- WIP, lots of ideas in the pilepine …

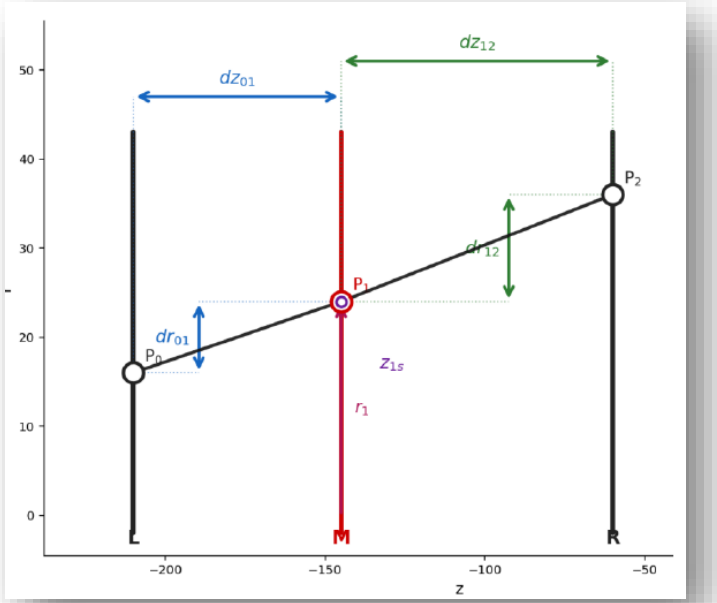

LHCb – Analysis Facilities with the HLT1 GPUs

General idea: exploit the LHCb trigger CPU and GPU resources as Analysis Facilities for offline analysis related work

- SWAN is an analysis facility, hence collaborate with the CERN-IT / SWAN team to leverage as much as possible – joint pilot project since 2025 :)

- **Proceed in steps, taking into account internal and external constraints**:
  - 8-GPU mini cluster set up and running since late Summer 2025
    - Independent of the special Online environment, validates the end-to-end chain, accessible via a SWAN QA private instance
  - Attach the Online GPUs as external SWAN resources during LS3

- Along the way, also got the **default LHCb software environment available in SWAN** (custom envs work)

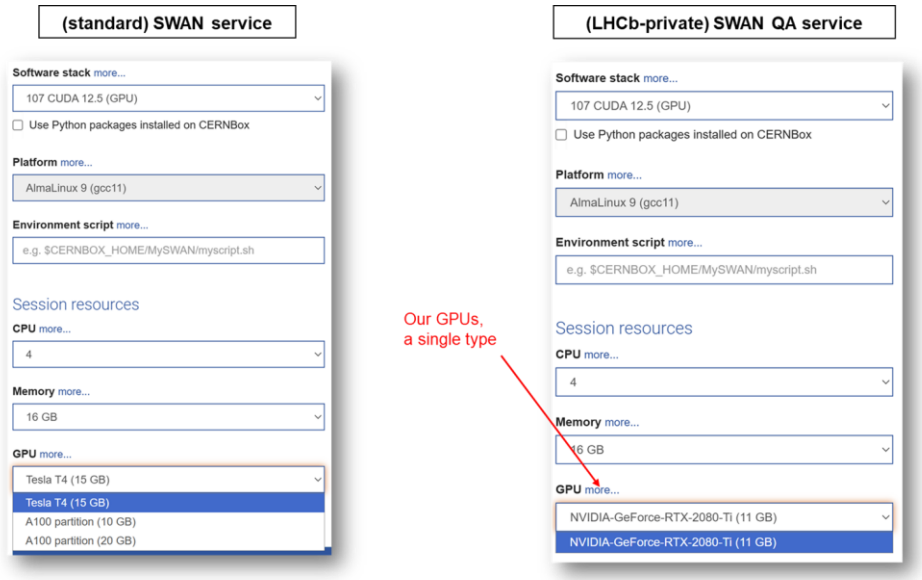

- CHEP 2026 talk jointly with CERN's SWAN team next week:

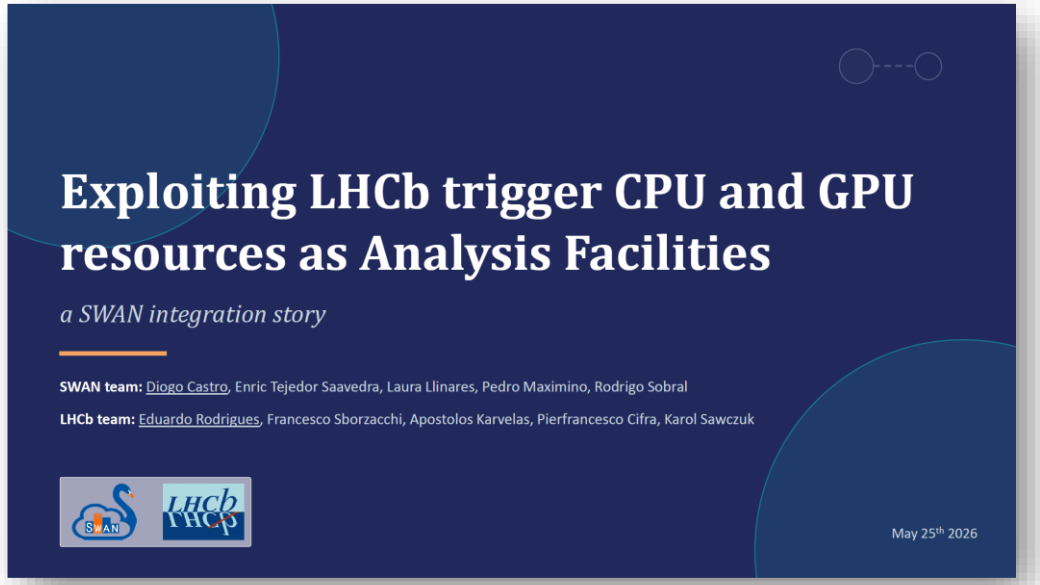

AI4HEP grant/initiative

Eduardo, Monica, Rob

- **Strategically very important grant for us, a step towards much more**
- Project started in Jan. 2026, £200k together with UCL, IC, Sheffield; Monica is a co-PI
- Several mini-projects funded via the grant. Liverpool runs a more substantial project (see next slide) and takes care of the website https://ai.hep.ac.uk/
- AI4HEP ran the 3rd Workshop on Enabling AI in HEP experiments and theory in London in March 2026
- 18 May, contribution from Monica to US Blueprint Workshop "Towards a National-Scale AI Collaboration in HEP", where AI for Particle Physics in the UK (hence AI4HEP) was presented

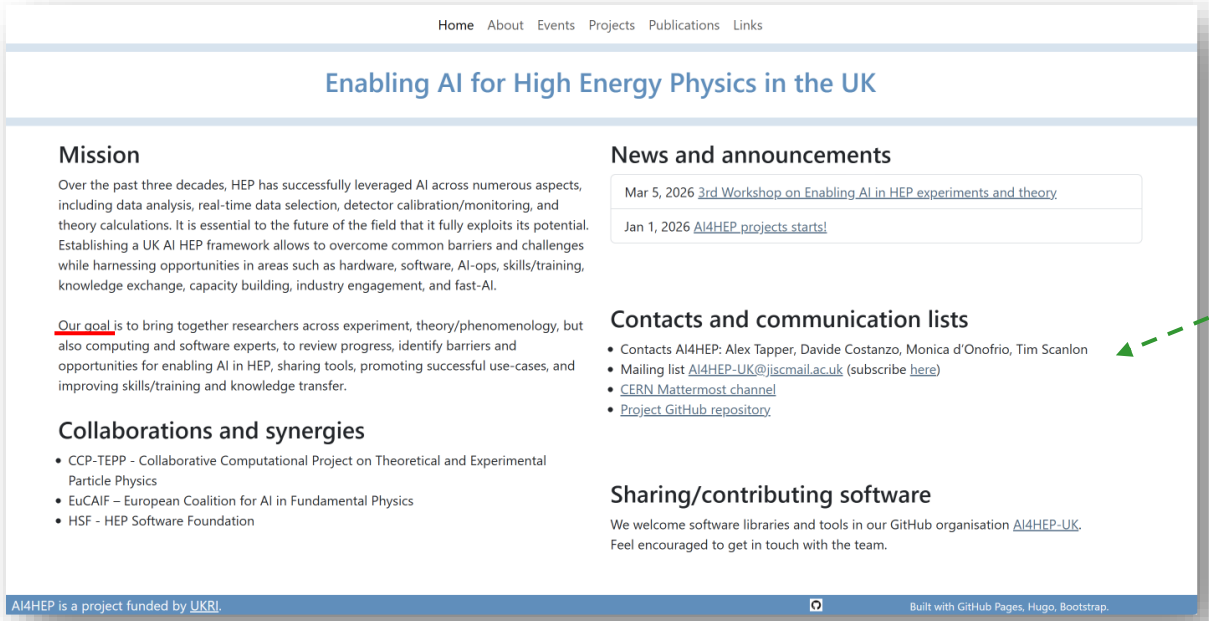

Feel encouraged to subscribe !

AI4HEP grant/initiative – Liverpool project

Eduardo, Monica, Rob

"Smart Monitoring: AI-Driven Analysis of Distributed Systems"

Can we move towards a more efficient usage of resources through "intelligent" monitoring? [see full talk]

Implement proof-of-concept focused on LHC-related computing operations

- Use local (small) AI models to detect and analyse issues through monitoring frameworks
- Leverage monitoring platform Zabbix connected to locally run AI model to query state of monitored systems and services (interactively or automatically)
  ➞ enable dynamic detection and analysis

Results so far (on a local computing site / Tier-2):

- Demonstrator model provides insight into state of system & faster responses
- Plans to expand with further interactive investigation and autonomous actions

Key Considerations for scaling up:

- Accuracy / insight, efficiency / responsiveness, safety in autonomous actions, maintainability / environmental sustainability
- If we get things right, we can reduce efforts and improve response with minimal environmental cost

SusTraIN grant/initiative

Eduardo, Monica, Rob

SusTraIN = Sustainable Training Initiative for NetDRIVE:

- NetDRIVE (UKRI Net-Zero Digital Research Infrastructure Vision and Expertise) Community Project
- Started in April 2026 jointly between Manchester (PI), Liverpool (Monica and Eduardo as co-PIs), Sheffield, the SSI (UK Software Sustainability Institute)
- Rather small grant but still important to be on the map on sustainable computing (+ training)
- **3 main topics towards green/sustainable computing**: Review of training and needs (WP1), training material (WP2), training catalogue (WP3)

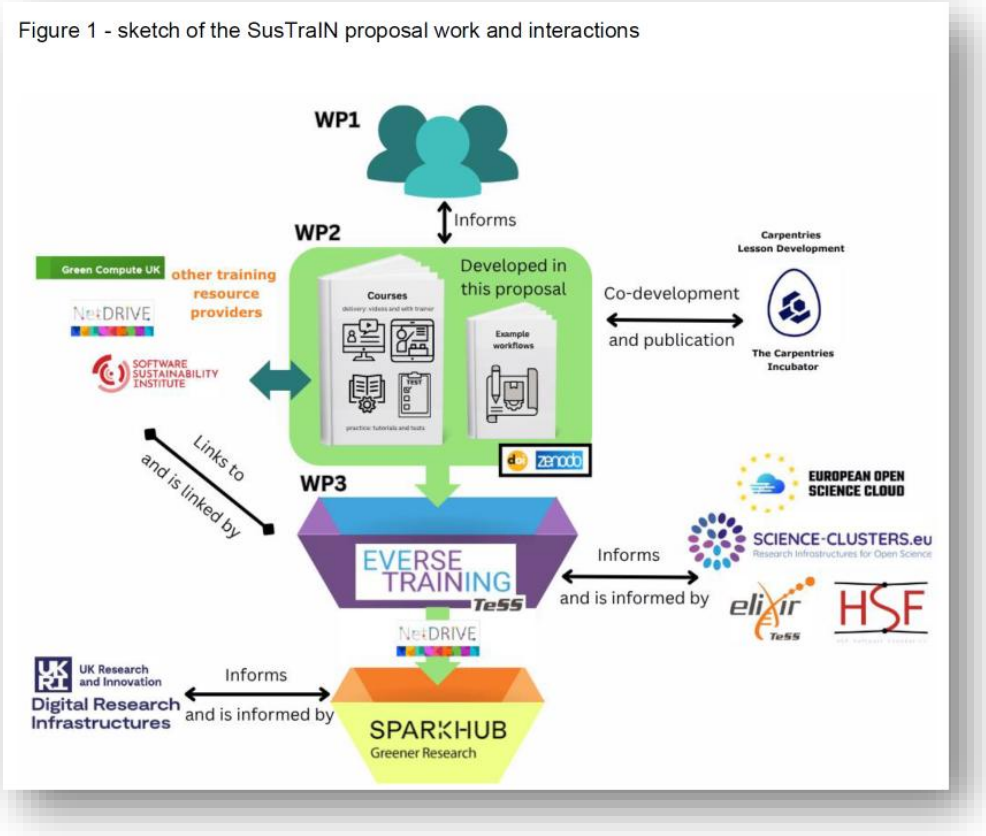

Liverpool contributions:

- **Liverpool** is responsible for WP2, **to develop short, hands-on training modules and reproducible workflow examples**, integrating CPU/GPU energy-profiling, FAIR metadata, and Zenodo-ready releases
- **Translate requirements** (particularly around AI/ML examples) **into executable training artefacts and student-tested materials**
- **Connect also with the HSF**

Community work – HEP Software Foundation (HSF) and related

Eduardo

https://hepsoftwarefoundation.org/

Roles:

- HEP Software Foundation (HSF) Steering Group chair, since June 2025
- Worldwide LHC Computing Grid (WLCG) Management Board software representative, 09/2023 – 08/2025
- Joint WLCG/HSF workshop committee programme member for 2025 and future 2026 editions

Talks:

- Invited talk "**HEP software: past, present, future**" at the Open Symposium on the European Strategy for Particle Physics, Palazzo del Casinò, Venice Lido (Italy), 23 - 27 June 2025
- EVERSE Community Engagement Event, CERN, Geneva (Switzerland), 5 February 2026
  HSF - The HEP Software Foundation
- XVII CPAN Days, ADEIT, Valencia (Spain), 19 - 21 November 2025
  HSF - The HEP Software Foundation

Scikit-HEP project

"Scikit-HEP is an open-source community-driven and community-oriented project with the goal of providing an ecosystem for particle physics data analysis in Python, fully integrated with the wider scientific Python ecosystem.

It expands the typical Python data analysis tools for particle physicists with packages spanning the spectrum from general scientific libraries for data manipulation to domain-specific libraries."

- **Recent work** focused on consolidating the project – it is turning 10 later in the year!
- Continue maintaining several libraries
- Project libraries continue to be cited, see here

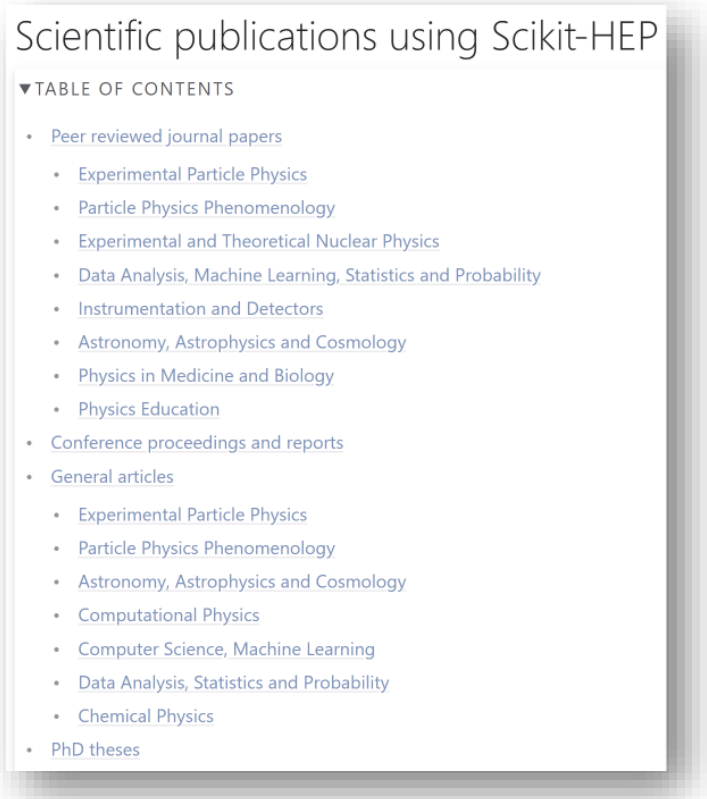

Talks:

- **The Scikit-HEP project,** Future of SHiP computing workshop, CERN, February 2026
- **The Scikit-HEP project - news and future directions,** WLCG/HSF Workshop 2025, IJCLab, May 2025

Seminars:

- **The Scikit-HEP project: overview and future**, PUNCHLunch Seminar, Germany, June 2025

Other collaborations

Healthcare (*Monica, Nikos, Rob*):

- **CNN and Graph Transformer-based pipelines for healthcare applications** (predictions of eye disease evolution), with HLS team from Yalin Zheng
- Also looking at the possibility to develop further **collaborations with Engineering** (Dr Yifan Zhou, in the group of Simon Maskell) to work on **AI and xAI pipelines for healthcare applications**, and get further support under the Civic Health Innovation Liverpool (CHI-L) centre

Space Engineering & ESA (*Monica*) :

- **Collaboration with Space engineers at Liverpool and ESA** through PhD student Amelia Samuel: https://www.liverpool.ac.uk/doctoral-training/innovation-in-data-intensive-science/meet-our-students/amelia-samuel/
- Monica is 2nd supervisor of project with Dr Stefania Soldini from Aerospace Engineering Department in the School of Engineering. Amelia is also part of the LIV.INNO programme

Public relations and outreach

- Outcome of a *panel discussion at an Economist Fusion Fest event*
  - Tara was co-chair
- **"Challenges and opportunities for AI to help deliver fusion energy"** (https://arxiv.org/abs/2603.25777)

Abstract. There is great potential for the application of AI tools in fusion research, and substantial worldwide benefit if fusion power is realised. However, using AI comes with its own challenges, many of which can be mitigated if responsible and robust methodologies are built into existing approaches. To do that requires close, long-term collaborations between fusion domain experts and AI developers and awareness of the fact that not all problems in fusion research are best tackled with AI tools. In April 2025, experts from academia, industry, UKAEA and STFC discussed how AI can be used to advance R&D in fusion energy at the first edition of The Economist *FusionFest* event. This Perspective is an expanded and updated summary of the round table discussion, providing more context and examples.

ARTIFICIAL INTELLIGENCE | FEATURE

How AI can help (and hopefully not hinder) physics

29 Jul 2025

The use of artificial intelligence is already becoming commonplace in physics, but could physics also help AI? **Tara Shears** examines the relationship between the two fields following a survey and report by the Institute of Physics

- **Physicsworld** article [link]

Publications

Eduardo – on LHCb, HSF/community, Scikit-HEP:

- **The LHCb Stripping Project: Sustainable Legacy Data Processing for High-Energy Physics**
  N. Grieser et al., Comput Softw Big Sci 9, 21 (2025); doi:10.1007/s41781-025-00151-6
- **The LHCb Sprucing and Analysis Productions**
  A. Abdelmotteleb et al., Comput Softw Big Sci 9, 15 (2025); doi:10.1007/s41781-025-00144-5

← **LHCb work**

- **Vector: JIT-compilable mathematical manipulations of ragged Lorentz vectors**
  Saransh Chopra et al., Journal of Open Source Software, 10 (109), 7791; doi:10.21105/joss.07791

← **Scikit-HEP work**

- **The Critical Importance of Software for HEP**
  Christina Agapopoulou et al., HSF prepared inputs for the European Particle Physics Strategy Update
  Eur. Phys. J. C 85 (2025) 1142; doi:10.1140/epjc/s10052-025-14571-6, arXiv:2504.01050 [hep-ex]

← **HSF work**

Marina Maneyro, Samuel Godwood – on Quantum Computing:

- **Fault-Tolerant Resource Comparison of Qudit and Qubit Encodings for Diagonal Quadratic Operators**
  https://arxiv.org/abs/2604.26792

Little poll – usage of GenAI

Hands up who is using Generative AI in one form or another for their work?

- Interesting to assess the percentage of colleagues
- (We won't have time in this talk to separate academics from PhD students + early career)

Who is paying subscriptions on their own?

- Corollary: should one be sending a message higher up on providing subscriptions for free to people …?

Thank you for listening !