

Blueprint Workshop: Towards a National-Scale AI Collaboration in HEP Closeout Summary

Organizers: Peter Elmer (Princeton University), Tulika Bose (U.Wisconsin-Madison),
Verena Martinez Outschoorn (U.Massachusetts-Amherst), Ian Fisk (Flatiron Institute)



This workshop is being organized by the Coordinating Panel for Software and Computing (CPSC) which is hosted by the Division of Particles and Fields (DPF) of the American Physical Society (APS).

Support for this workshop has been provided by the Institute for Research and Innovation in High Energy Physics (IRIS-HEP) through NSF award PHY-2323298 and by the FAIROS-HEP project through NSF awards grants OAC-2226378, OAC-2226379 and OAC-2226380.



“Building an AI-Native Research Ecosystem” Whitepaper

arXiv > hep-ex > arXiv:2602.17582

Search...

All fields

Search

Help | Advanced Search

High Energy Physics – Experiment

[Submitted on 19 Feb 2026]

Building an AI-native Research Ecosystem for Experimental Particle Physics: A Community Vision

Thea Klæboe Aarrestad, Alaa Abdelhamid, Haider Abidi, Jahred Adelman, Jennifer Adelman-McCarthy, Shuchin Aeron, Garvita Agarwal, Usman Ali, Cristiano Alpigiani, Omar Alterkait, Mohamed Aly, Oz Amram, Saeed Ansari Fard, Aram Apyan, John Arrington, Marvin Ascencio-Sosa, Mohammad Atif, Aneesha Avasthi, Muhammad Bilal Azam, Bhim Bam, Joshua Barrow, Rainer Bartoldus, Amit Bashyal, Aashwin Basnet, Ayse Bat, Lothar A. T. Bauerdick, John Beacom, Chris Bee, Michael Begel, Matthew Bellis, Rene Bellwied, Rakitha Beminiwattha, Gabriele Benelli, Douglas Benjamin, Catrin Bernius, Binod Bhandari, Avinay Bhat, Meghna Bhattacharya, Saptaparna Bhattacharya, Prajita Bhattarai, Sudip Bhattarai, Wahid Bhimji, Jianming Bian, Burak Bilki, Mary Bishai, Kevin Black, Kenneth Bloom, Brian Bockelman, Johan Sebastian Bonilla Castro, Tulika Bose, Nilay Bostan, Othmane Bouhali, Dimitri Bourilkov, Dominic Brailsford, Gustaaf Brooijmans, Elizabeth Brost, Maria Brigida Brunetti, Quentin Buat, Brendon Bullard, Jackson Burzynski, Paolo Calafiura, Rodolfo Capdevilla, Fabian Andres Castaño Usuga, Raquel Castillo Fernandez, Fabio Catalano, Viviana Cavaliere, Flavio Cavanna, Giuseppe Cerati, Aidan Chambers, Maria Chamizo-Llatas, Philip Chang, Andrew Chappell, Arghya Chattopadhyay, Sergei Chekanov, Jian-ping Chen, Yi Chen, Zhengyang Chen, J. Taylor Childers, Hector Chinchay, Yuan-Tang Chou, Tasnuva Chowdhury, Neil Christensen, Wonyong Chung, Rafael Coelho Lopes de Sa, Simon Corrodi, Kyle Cranmer, Matteo Cremonesi, Roy Cruz, Mate Csanad, Mariarosaria D'Alfonso, Carlo Dallapiccola, Daine Danielson, Sridhara Dasu, Gavin Davies, Kaushik De, Patrick de Perio, Klaus Dehmelt, Marco Del Tutto, Carlos Ruben Dell'Aquila, Sarah Demers et al. (359 additional authors not shown)

Experimental particle physics seeks to understand the universe by probing its fundamental particles and forces and exploring how they govern the large-scale processes that shape cosmic evolution. This whitepaper presents a vision for how Artificial Intelligence (AI) can accelerate discovery in this field. We outline grand challenges that must be addressed to enable transformative breakthroughs and describe how current and planned experimental facilities can implement this vision to advance our understanding of the vast and complex physical world from the smallest to the largest scales. We show how facilities currently under construction, such as the HL-LHC, DUNE and soon EIC, can both benefit from and serve as proving grounds for this vision, while also enabling a longer-term goal for how future experiments -- like FCC-ee at CERN, IceCube-Gen2, a Muon Collider in the U.S., and smaller to mid-scale projects -- can be fully AI-native. We describe how a truly national-scale collaboration, jointly managed across large funding partners, and involving both DOE laboratories and universities, can make this happen.

<https://arxiv.org/abs/2602.17582>

Access Paper:

[View PDF](#)
[HTML \(experimental\)](#)
[TeX Source](#)

 [view license](#)

Current browse context:

hep-ex

[< prev](#) | [next >](#)

[new](#) | [recent](#) | [2026-02](#)

References & Citations

[INSPIRE HEP](#)
[NASA ADS](#)
[Google Scholar](#)
[Semantic Scholar](#)

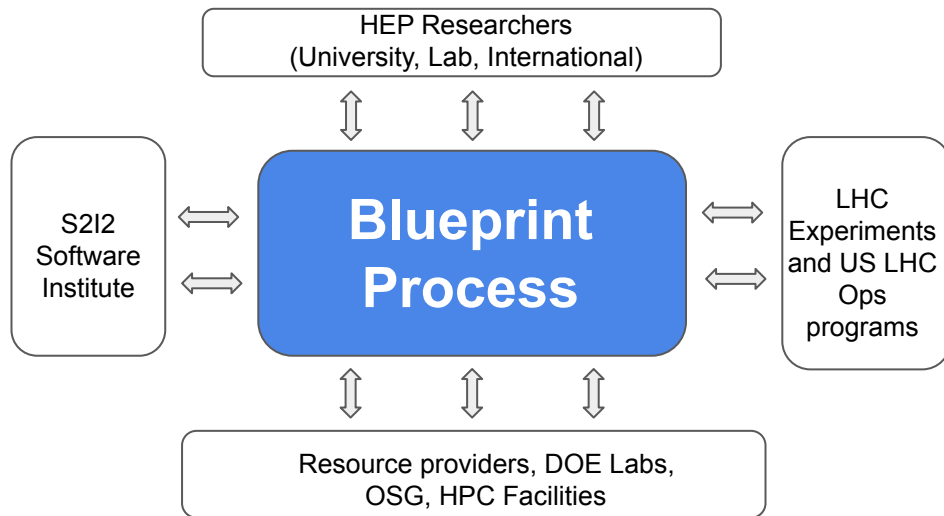
Export BibTeX Citation

Bookmark



Commitment to Joint DOE and NSF Blueprint Activity

- Drive the evolution of R&D efforts to address the software & computing challenges of the HL LHC, co-sponsored by:
 - **US LHC Ops program**
 - **S2I2**
 - **OSG**
 - **CCE**
- Involving the DOE facilities, and key personnel at both DOE labs and US Universities.
- Long term sustained set of workshops to drive coherence across projects and experiments.



Monday 18 May

Short (<10 minute) presentations.

Very short big picture overview, with most of the talk hopefully contributing directly to the (AI Collaboration) discussion.

There is a google doc to make notes about emerging themes.

At the end of the day we will discuss identified themes for follow-on discussion.


Workshop Dinner at 19:00 for those attending in person.

MONDAY 18 MAY		
12:00 → 13:00	Lunch	1h
13:00 → 13:20	Introduction and Goals for the Workshop	20m
13:20 → 14:00	Initial Discussion	40m
14:00 → 14:10	TREASURE Project Speaker: Haider Abidi (BNL)	10m
14:15 → 14:25	Next Generation DQM Project Speaker: Walter Hopkins (Argonne National Laboratory)	10m
14:30 → 14:40	Open Data and AI Speaker: Zach Marshall (Lawrence Berkeley National Laboratory)	10m
14:45 → 14:55	American Science Cloud and Demonstrators Speaker: Taylor Childers (Argonne Nat. Lab.)	10m
15:00 → 15:30	Coffee Break	30m
15:30 → 15:40	Knowledge Extraction Project Speaker: Nesar Ramachandra (Argonne National Laboratory)	10m
15:45 → 15:55	CERN Strategy for AI and Organisational Plans Speaker: Maurizio Pierini (CERN)	10m
16:00 → 16:10	INFN Speaker: Daniele Bonacorsi (University of Bologna / INFN)	10m
16:15 → 16:25	UK Speakers: Claire Shepherd-Themistocleous (Rutherford Appleton Laboratory, UK), Davide Costanzo (University of Sheffield (GB))	10m
16:30 → 16:40	Neutrinos Speaker: Alexander Himmel (Fermi National Accelerator Laboratory)	10m
16:45 → 16:55	Thoughts on Smaller Experiments Speaker: Lindley Winslow (MIT)	10m
17:00 → 17:15	Organization Discussion	15m
19:00 → 21:30	Workshop Dinner	2h 30m

Four Grand Challenges Around the Experimental Lifecycle

“Building an AI-native Research Ecosystem” - Focuses on potential system-level change to the research ecosystem.

“Building an AI-Native Research Ecosystem” - Inclusive of path from R&d to r&D to eventual deployment include scale-up and infrastructure connections.




National Scale Collaboration?

U.S. Institutions (**blue dots = universities**, **red dots = DOE labs**) that are currently part of large national-scale funded projects:

- U.S. ATLAS Operations (DOE/NSF, 33 institutions)
- U.S. CMS Operations (DOE/NSF, 44 institutions)
- DUNE Operations (DOE, 16 institutions)
- HEP-CCE (DOE, 4 labs)
- IRIS-HEP (NSF, 16 universities)
- A3D3 (NSF, 12 universities)

Additional universities that are part of the U.S. ALICE, U.S. ATLAS, U.S. CMS and (U.S.) DUNE scientific collaborations are shown in green.

Taken together, these overlapping experiments and projects tightly connect 9 DOE labs and 124 different universities and colleges.



A central focus of the workshop will be discussions on how such a national-scale collaboration could be effectively bootstrapped and organized, as well as identifying priorities over the near and medium term. Such a collaboration, including both universities and DOE laboratories, has a number of potential benefits and opportunities, including:

- a forum for developing a shared roadmap toward an AI-native research ecosystem
- a platform for enabling larger-scale efforts while integrating expertise and resources from across the community
- clear pathways to scale and transition promising R&D activities towards deployment within experiments
- a structure to develop and provide the services needed for the experiments to fully leverage evolving national infrastructure (e.g., AmSC, NAIRR) and large-scale data resources
- a means for fostering both collaboration and healthy competition through the development of shared benchmarks and the development and curation of the supporting datasets
- accessible entry points for individual university groups to engage and contribute
- a coherent framework for cross-experiment activities, enabling solutions with impact to deliver benefit to multiple experiments
- a clear structure for US engagement with international efforts, including those at CERN, and other national and regional collaborations
- a community structure for engagement with industry
- a structure for large scale workforce development activities leveraging both lab and university expertise

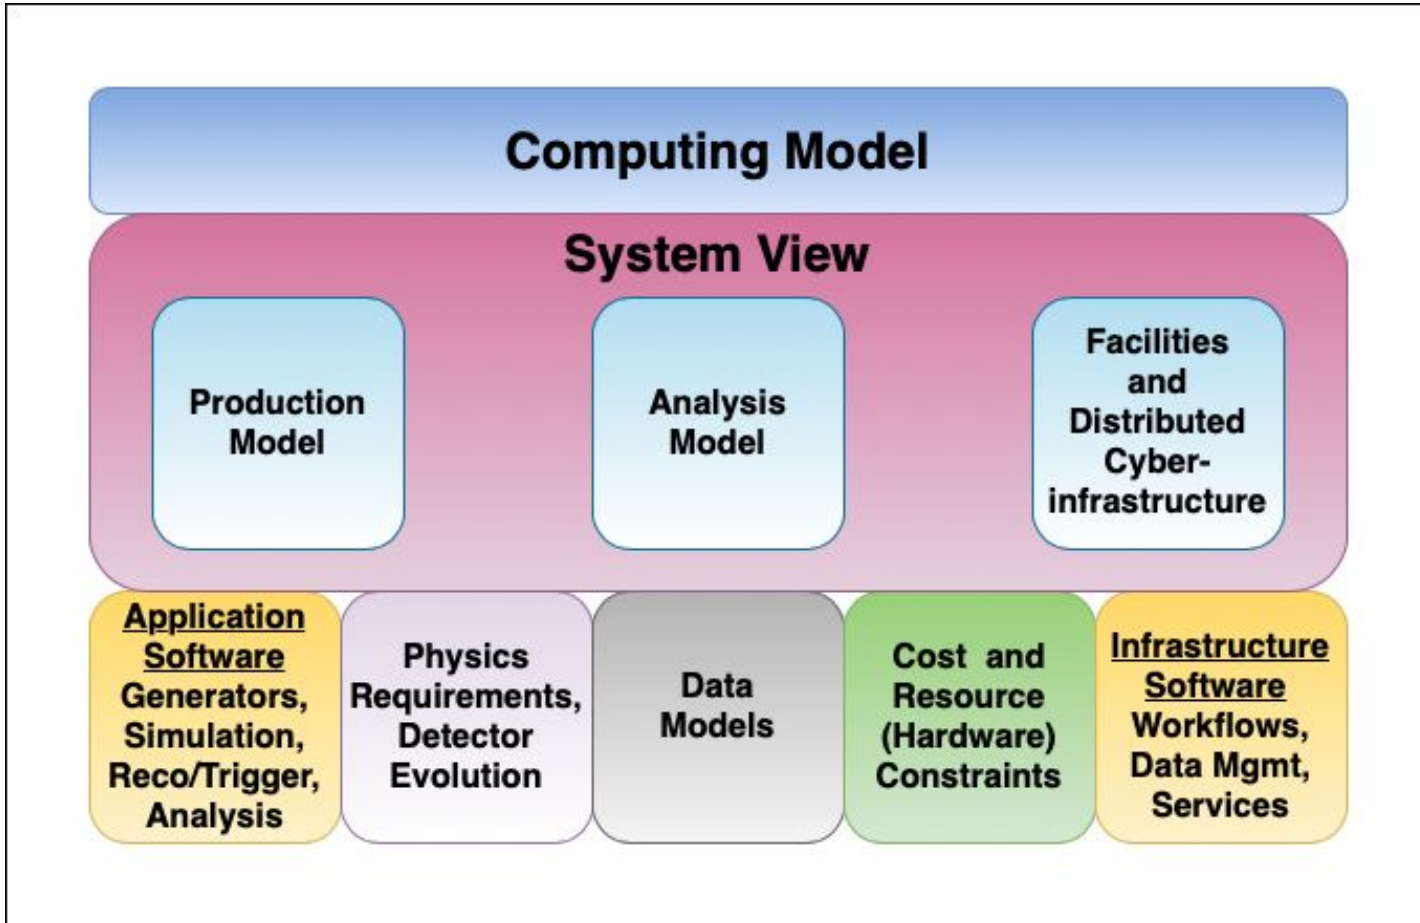
This first workshop will consider the potential impact of these opportunities, identifying near-term priorities and exploring means to enable additional initiatives as resources become available. We expect that this discussion will continue in subsequent community workshops during 2026.

A central focus of the workshop will be discussions on how such a national-scale collaboration could be effectively bootstrapped and organized, as well as identifying priorities over the near and medium term. Such a collaboration, including both universities and DOE laboratories, has a number of potential benefits and opportunities, including:

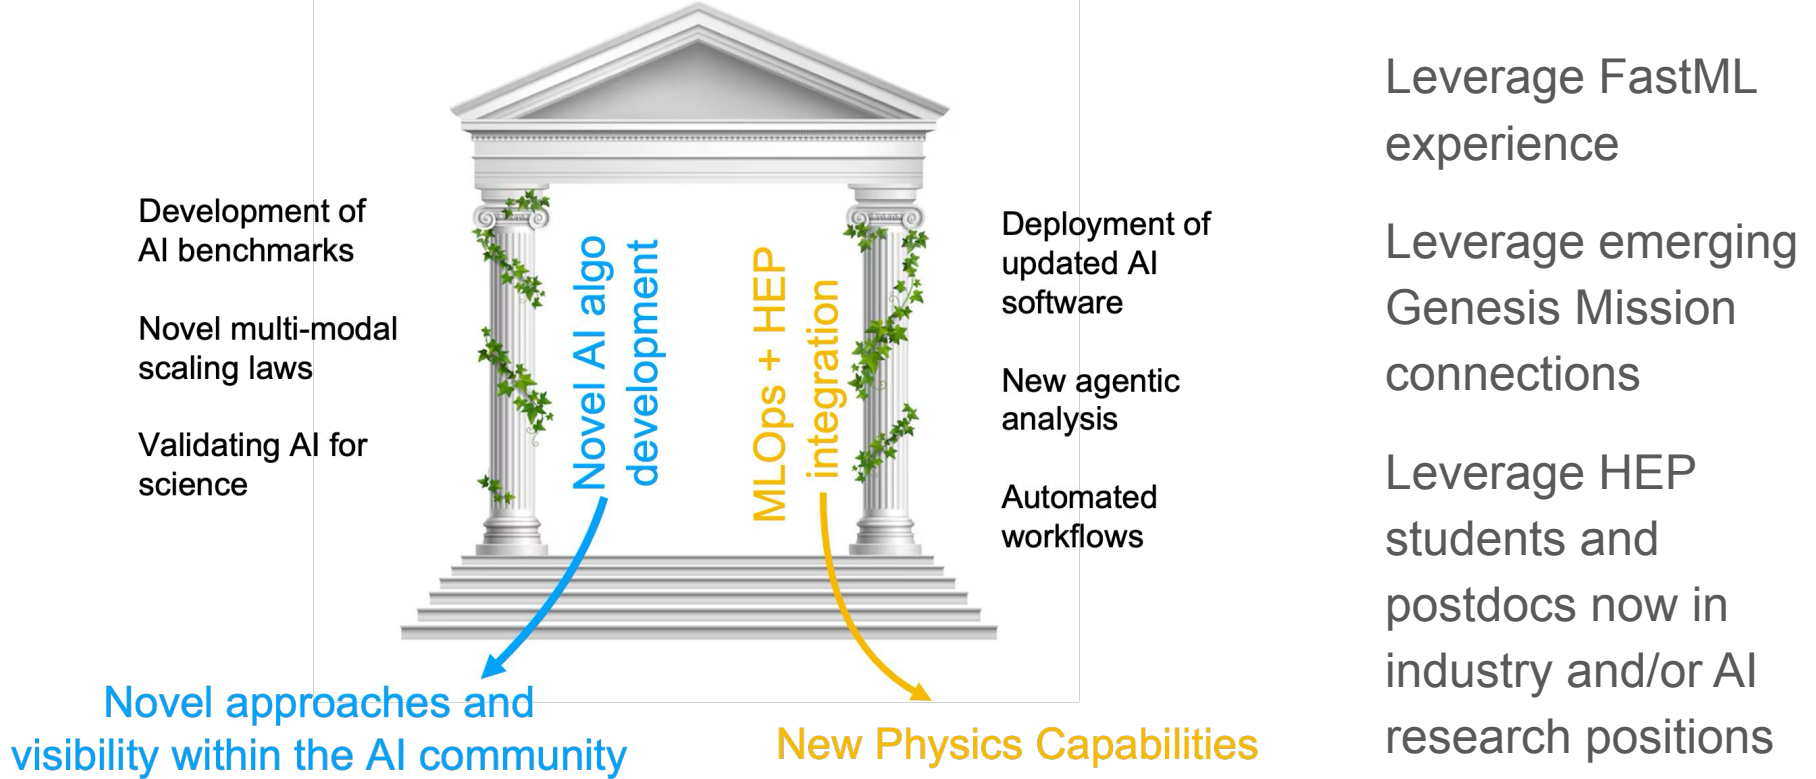
- a forum for developing a shared roadmap toward an AI-native research ecosystem
- a platform for enabling larger-scale efforts while integrating expertise and resources from across the community
- clear pathways to scale and transition promising R&D activities towards deployment within experiments
- a structure to develop and provide the services needed for the experiments to fully leverage evolving national infrastructure (e.g., AmSC, NAIRR) and large-scale data resources
- a means for fostering both collaboration and healthy competition through the development of shared benchmarks and the development and curation of the supporting datasets
- accessible entry points for individual university groups to engage and contribute
- a coherent framework for cross-experiment activities, enabling solutions with impact to deliver benefit to multiple experiments
- a clear structure for US engagement with international efforts, including those at CERN, and other national and regional collaborations
- a community structure for engagement with industry
- a structure for large scale workforce development activities leveraging both lab and university expertise

This first workshop will consider the potential impact of these opportunities, identifying near-term priorities and exploring means to enable additional initiatives as resources become available. We expect that this discussion will continue in subsequent community workshops during 2026.

Software, System View & Computing Models



A natively “AI-for-Physics/Physics-for-AI” collaboration



“Proto-collaboration” Near Term Bootstrap Activities

Engage with existing funded projects (and emerging funded projects in the rest of FY26) and create links between them and to the community

Identify a set of priority areas of interest (see subsequent slides)

Identify a set of activities and/or demonstrators within those priority areas that are doable on the time scale of the next six months

Being to create environments to catalyze both working collaborations and connections for eventual proposals

Checkpoint and engage community regarding the proto-collaboration at APS DPF @FNAL in July (parallel session)

Aim for a first “collaboration” meeting in the fall.

Continue discussion about links between this entity and international partners

Topics of discussion

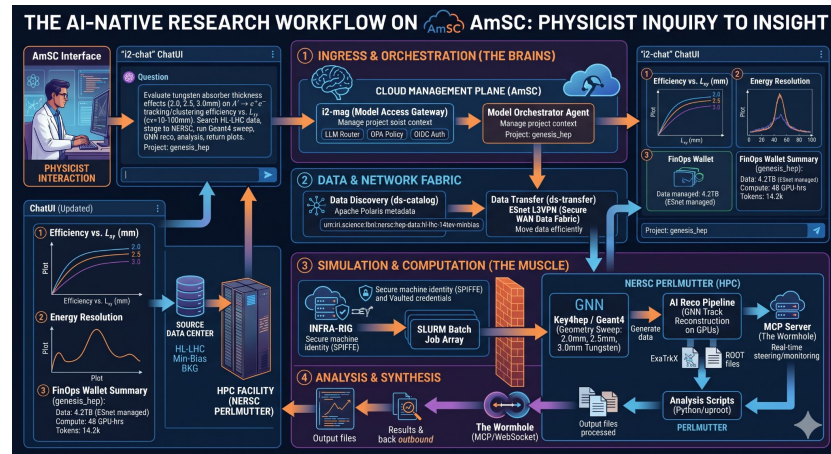
1. American Science Cloud
2. TREASURE
3. AI Operations In Proto-Collaborations
4. Small/Medium Scale Experiments

Breakouts

5. Computing Resource Strategy in the age of AI
6. Agentic AI
7. Foundation Models
8. Open Data & Benchmarks
9. Workforce Development

Genesis Mission | AMERICAN SCIENCE CLOUD

- Model Access Gateway (think OpenAI API) with industry model access + Federated ID by **July** to support Genesis.
- Minimal Viable Product by **October** adds ChatUI, Data Catalog & Transfer/Streaming, AI experiment tracking, Agent/MCP Catalog. Federated ID for inside/outside DOE.
- Equinox (ALCF), Lux (OLCF) each 10k GPU systems expected online before end of FY26.
- Solstice (ALCF) is a 100k GPU system in next two years.
- In addition, the LHC experiments were “invited” to become “Infrastructure Partners” for the emerging Genesis Platform
- Fermilab Data Platform is an Infrastructure Partner.



TREASURE

- Transform heterogeneous collider data into standardized, tokenized, AI-ready representations that enable foundation models to learn across experiments, detectors, and data-taking eras.
- **Next goals and milestones:**
 - Establish Treasure workflow on AmSC & provide data through ModCon
 - Create a context-aware event-level tokenization
 - Focus on low-level information tokenization
- Will provide support for Genesis Phase 1 proposals or any other collaborators
 - Tokenized data, libraries, workflow are there as basis for different **Foundation Models**
 - Depending on how many genesis proposals/interested people, will setup a common forum to share ideas and solve problems
- **Aim to bring together Phase 1 + collaborators for a combined submission for Phase 2 in December**

the TREASURE Pipeline



AI for operations within proto-collaboration

- Continue to have open ~weekly meetings and connect with other operations efforts (perhaps serve as a forum for shared experiences and challenges)
- Already is a framework for cross-experiment activities, enabling solutions with impact to deliver benefit to multiple experiments, which can be expanded
- Currently have hackathon/discussions every ~2 months at Argonne/Fermilab. Maybe we could expand this?
- Serve as an accessible entry point for individual university groups to engage and contribute

Small/Medium Experiments

- Diverse physics goals but share common challenge in integrating low level and high level data to effectively harness AI into all aspects of the experimental life cycle.
 - In general, do not have a forum to bring together the community.
- Resource acquisition is hard for computationally intensive projects.
- Expertise is diffuse across efforts.
- Idea:
 - establish connections with main effort (LHC/Dune): goal is to stay connected to major operations, not to do things on our own.
 - small/medium experiment forum to collect commonalities/needs; pool resources to transfer/implement commons solutions. Might differ across communities.
 - potential action plan:
 - Forum in June/July (DPF pre-event?)
 - 1-2 proof of principle common solutions prototyped for 6 months meeting

Computing Resource Strategy in the AI Era

AI is reshaping the computing landscape: the rapid build-up of AI infrastructure is changing hardware markets, facility priorities, and access models for scientific computing.

The traditional WLCG model will need to evolve: a primarily dedicated, CPU-based infrastructure may not remain sustainable or sufficient as HPC, AI facilities, accelerators, and externally allocated resources become increasingly central.

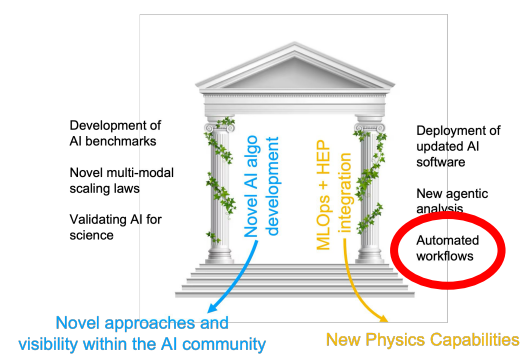
HEP needs a coordinated community response: a workshop and white paper could define the future computing model, including the role of AI workflows, HPC access, inference facilities, and on-prem resources.

Key questions require urgent assessment: the scale of training, tuning, and inference workflows; how these augment or replace existing workloads; and the resulting implications for networks, data access, storage, and scheduling.

Engagement with HPC and AI facilities must be at the community level: HEP should pursue common interfaces, common allocation mechanisms, shared development resources, and coordinated access models aimed at enabling abundance rather than scarcity.

Agentic AI

Automated agent driven workflows for Analysis, Operations, Coding, Hardware Design, etc.



How might this Collaboration Help?

- Holder of benchmarks and their open data and validation procedures.
- Agentic Tutorial repository
- Repository of standard agents, skills, etc. that are cross experiment
 - Example:
<https://github.com/usatlas/marketplace>
- Knowledge repository (e.g. all papers from ATLAS CMS ready for ingestion, ...)

In 6 months?

- Hackathon (problems already solved, build benchmarks & validation)
- Open Computing Facility
- Real Use (published, shifter replacement) deep dive.
- Operations: open data - towards open-data tutorials
- Lightning talks of progress (Genesis Funded?)

Ran out of time to discuss industry connections, did have thoughts on organization but didn't place

Foundation Models: How can our communities benefit from this collaboration?

Creates opportunities/mechanism for cross-experimental models/collaboration

Collectively advocate and organize efforts for foundation models:

- identify HEP-wide needs
- physics datasets
- advocate for collective computational resources
- Create suite of physics, ML, computational benchmarks
- Identify needs for physics-based guardrails / constraints / validation

Find inter-experimental connections surrounding specific topics: e.g. physics-native FM, encoders, sensor to symbolic models, online reconstruction

build from or **interface with existing efforts** such as TREASURE

Engage in important conversations: E.g. what does physics bring to broader AI/ML efforts? Non-”standard” data structures

Create or curate common platforms to store / share models, data, tools, etc

Foundation Models: What can we do in the next 6 months?

Technical:

define/provide downstream benchmarks for community use (physics, computational)

define foundation model needs: data structures to encode, scope of models

representation learning frameworks for analysis of transferability of models

Community:

share best practices, existing efforts, pitfalls, common goals

identify computational resources for large-scale FM training, look at potential for community-wide resources

What can we learn from existing efforts (industry, etc)?

Define list of scopes of interest to sub-communities

Open Data Breakout

White paper coming soon (“A vision for Open Data in the Age of AI”), target 1–2 months to a readable draft, covering:

- The current state of open data releases
- Existing uses and limitations (including hardware and software limitations)
- Opportunities (more data, other formats, educational opportunities)
- Requirements for expansion (metadata, provenance, hardware resources, machine-readable documentation, ...)
- Open questions about AmSC and what it will provide (authentication, duplication, etc) — aspirations and reality

Key points from discussion (initial 6-month targets **in bold**):

- We should definitely try to get a **hosted set of open data on some cloud provider**
- Documentation needs to be machine-readable (or Skill readable)
- Integrated systems are great (hardware, software stacks, model serving, intermediate data set storage, citation services)
- Should make sure **other communities are well integrated into the vision** (neutrino and eIC in particular)
- Need to think more about workflow capture and agent histories for documentation / ‘publication’ (maybe after embargo?)
- Should integrate support for benchmarks and common software solution development
- Need to think about harmonized monitoring (serving data from multiple sources, need to track who’s using what in a coherent way, not issuing multiple DOIs for identical data...)
- Should watch experiment policies — harmonization among big experiments, inclusion of small experiments, relatively low barriers so that we get *something* out

Blueprint Workshop: Workforce Development Summary

- **Audience**: start with graduate students, as we have generally defined benchmarks for graduate students
- **Fundamental Questioning**: is it possible to train a student to have intuition about, for example, understanding plots that are made with agentic AI?
- **Supporting Skills for Students**: what skills do we want our students to gain? Need to enable students to use new tools while training them to have intuition for physics; if students become managers of AI agents, we need to train students to be *critical* managers
- **Validation**: will become more important if students use agentic AI for analysis
- **Best Practices/Guiding Document**: develop set of best practices for use of AI for physics, validation of results, ethical use of AI, taking responsibility for the agents you use, etc
- **Curriculum**: work together to develop physics curriculum to address these issues
- **Learning Tools**: develop starter kit with resources on how to get started, recommendations, etc; develop webpage for these resources
- **Labs**: engage with labs on how to work with students

Blueprint Workshop: Workforce Development Summary

What can we do in 6 months?

- **Best Practices:** engage with programs (eg IRIS-HEP) to develop best practices → module in HEP software training
 - Aaron Higuera, Larry Lee, Sudhir Malik, Nadja Strobbe, Giordon Stark
 - Example: <https://usatlas.github.io/af-docs/ai/>
- **Starter Kit for AI Newcomers:** develop starter kit with resources on how to get started, recommendations, best practices, etc → online resource & training session
 - Toyoko Orimoto, Nadja Strobbe, Craig Group, Larry Lee
 - Webpage: Eliza Broust
- **Access to Agentic Tools:** could this collaboration provide access to subscriptions (eg Claude Code)
 - Mike Hance, Verena Martinez Outschoorn, Kamal Benslama
- Volunteer to join the effort! We will all benefit

Connections to (medium/high energy) NP facilities

- Experiments at
 - JLab, BNL (future EIC)
 - Possibility of building demonstrators on operating experiments in those facilities
- We are tackling similar problems to HEP experiments/facilities
 - Data quality monitoring → Hydra (computer vision, other techniques for DQM)
 - Detector design → AID2E
 - RL-based control for accelerators, polarized target systems, and coupling experiment+accelerator controls
 - HPDF + NP use case → Strategic LDRD in progress
 - Continual learning, optimization for accelerators (SNS, JLab, SLAC)
 - Online alignment, calibration and control of detectors
 - Foundation Models developed for multiple tasks (e.g., high-fidelity fast-sim, reco & analysis) for multi sub-detector systems
 - Fast ML for Streaming Readout
 - Agentic workflows for control & analysis

Collaboration with (medium/high energy) NP facilities

- Genesis Mission involvement
 - NARAD: Nuclear AI-Ready Accelerator Data
 - MOAT: Multi-Office Accelerator Team
 - SUF: Scientific User Facilities
 - + many more phase I & II proposals
- JLab Data and Computational Sciences directorate
 - Chief Data Office
 - Data Science Department
 - Performant Computing and Experimental Physics Software and Computing Infrastructure Groups
 - HPDF operations
 - Data Center Operations
- User community
 - AI4EIC (workshops, hackathons, tutorials)
 -

We are very interested in establishing *and maintaining* strong collaborations!

International Connections

CERN, INFN, UK participation in this meeting.

CERN is currently organizing a set of structured AI-RD collaborations, styled on the previous DRD structure.

INFN

- Post-workshop: start bilateral contact with INFN mgmt on possible resource allocation mechanisms, as part of a general discussion (refreshing WLCG?)
- Possible thematical areas of collaboration: see last 2 Bonacorsi's slides

UK

- Interested in collaborating, with a scope to be defined in a few months time
Priority area for the UK
- Funding needs to have potential industrial applications