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

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 Klaeboe 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.

Access Paper:

- View PDF
- HTML (experimental)
- TeX Source

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

https://arxiv.org/abs/2602.17582

Extra meeting links:

This workshop's Indico: https://indico.flatironinstitute.org/event/4120/

AI-native Ecosystem whitepaper: https://arxiv.org/abs/2602.17582

Slack channel (in IRIS-HEP Workspace):

#ai-collab-blueprint-may2026

Invite link for IRIS-HEP Workspace:

https://join.slack.com/t/iris-hep/shared_invite/zt-3yajggdtd-~yMOK1LfhkpehXIH_0h0kg

A google doc for taking notes will be circulated in the Slack channel.

Also other more targeted community efforts, for example:

arXiv > physics > arXiv:2602.22248

Search...

Help | Advan

Physics > Instrumentation and Detectors

[Submitted on 24 Feb 2026 (v1), last revised 10 Mar 2026 (this version, v2)]

Machine Learning on Heterogeneous, Edge, and Quantum Hardware for Particle Physics (ML-HEQUPP)

Julia Gonski, Jenni Ott, Shiva Abbaszadeh, Sagar Addepalli, Matteo Cremonesi, Jennet Dickinson, Giuseppe Di Guglielmo, Erdem Yigit Ertorer, Lindsey Gray, Ryan Herbst, Christian Herwig, Tae Min Hong, Benedikt Maier, Maryam Bayat Makou, David Miller, Mark S. Neubauer, Cristián Peña, Dylan Rankin, Seon-Hee (Sunny)Seo, Giordon Stark, Alexander Tapper, Audrey Corbeil Therrien, Ioannis Xiotidis, Keisuke Yoshihara, G Abarajithan, Sagar Addepalli, Nural Akchurin, Carlos Argüelles, Saptaparna Bhattacharya, Lorenzo Borella, Christian Boutan, Tom Braine, James Brau, Martin Breidenbach, Antonio Chahine, Talal Ahmed Chowdhury, Yuan-Tang Chou, Seokju Chung, Alberto Coppi, Mariarosaria D'Alfonso, Abhilasha Dave, Chance Desmet, Angela Di Fulvio, Karri DiPetrillo, Javier Duarte, Auralee Edelen, Jan Eysermans, Yongbin Feng, Emmett Forrestel, Dolores Garcia, Loredana Gastaldo, Julián García Pardiñas, Lino Gerlach, Loukas Gouskos, Katya Govorkova, Carl Grace, Christopher Grant, Philip Harris, Ciaran Hasnip, Timon Heim, Abraham Holtermann, Tae Min Hong, Gian Michele Innocenti, Koji Ishidoshiro, Miaochen Jin, Jyothisraj Johnson, Stephen Jones, Andreas Jung, Georgia Karagiorgi, Ryan Kastner, Nicholas Kamp, Doojin Kim, Kyoungchul Kong, Katie Kudela, Jelena Lalic, Bo-Cheng Lai, Yun-Tsung Lai, Tommy Lam, Jeffrey Lazar, Aobo Li, Zepeng Li, Haoyun Liu, Vladimir Lončar, Luca Macchiarulo, Christopher Madrid, Benedikt Maier, Zhenghua Ma, Prashansa Mukim, Mark S. Neubauer, Victoria Nguyen, Sungbin Oh, Isobel Ojalvo, Hideyoshi Ozaki, Simone Pagan Griso, Myeonghun Park, Christoph Paus, Santosh Parajuli, Benjamin Parpillon, Sara Pozzi, Ema Puljak et al. (22 additional authors not shown)

The next generation of particle physics experiments will face a new era of challenges in data acquisition, due to unprecedented data rates and volumes along with extreme environments and operational constraints. Harnessing this data for scientific discovery demands real-time inference and decision-making, intelligent data reduction, and efficient processing architectures beyond current capabilities. Crucial to the success of this experimental paradigm are several emerging technologies, such as artificial intelligence and machine learning (AI/ML), silicon microelectronics, and the advent of quantum algorithms and processing. Their intersection includes areas of research such as low-power and low-latency devices for edge computing, heterogeneous accelerator systems, reconfigurable hardware, novel codesign and synthesis strategies, readout for cryogenic or high-radiation environments, and analog computing. This white paper presents a community-driven vision to identify and prioritize research and development opportunities in hardware-based ML systems and corresponding physics applications, contributing towards a successful transition to the new data frontier of fundamental science.

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.

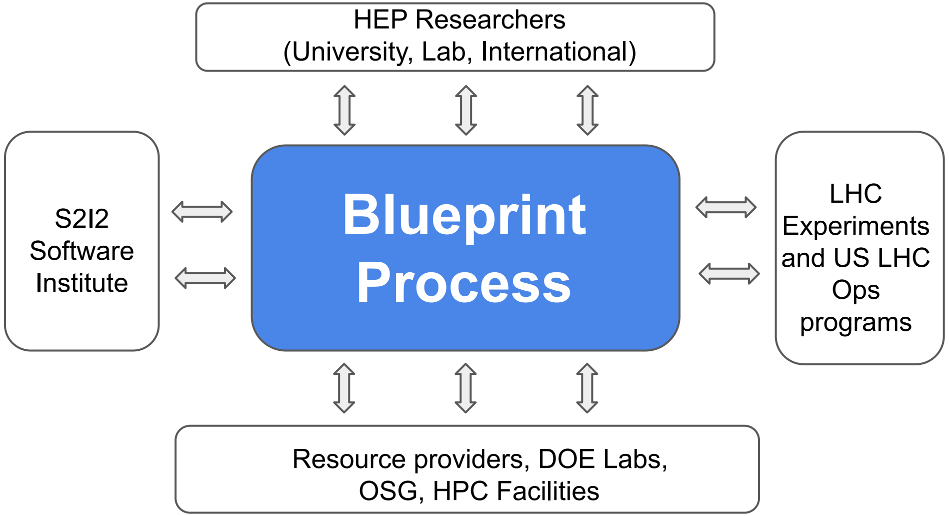

Blueprint workshops

Since 2017, we have carried forward this “blueprint” activity relative to HL-LHC Software and Computing R&D through the “Coordinated Ecosystem” workshops as well as a number of more targeted topical blueprint workshops.

The basic idea: we want to do more than just information exchange. We aim to have at least some actionable outcomes and/or improved organization as a result of the blueprint workshop.

Structure: less (and more targeted) presentations and plenty of time for discussion. Discussion time is not fully pre-planned and leaves space/time for potential ad-hoc parallel discussions. We aim to collect outcomes for a summary presentation at the end of the last day.

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

Tuesday 19 May

Breakfast for in-person participants at 8:30.

Three discussion blocks, we can break into multiple parallel sessions.

We will have a short discussion of the discussion plans at 9:00 and this will be summarized also for remote participants.

Additional zoom connections will be made available as needed.

Breakouts should produce short (1-2 slide summaries) which we will integrate into a single presentation for the closeout session at 15:15.

Tuesday 19 May

08:30	→ 09:00	Breakfast	30m
09:00	→ 10:30	Discussion	1h 30m
10:30	→ 11:00	Coffee Break	30m
11:00	→ 12:30	Discussion	1h 30m
12:30	→ 13:30	Lunch	1h
13:30	→ 15:00	Discussion	1h 30m
15:00	→ 15:15	Coffee Break	15m
15:15	→ 16:00	Closeout Session	45m

Several spaces available for discussions: plenary room, side meeting rooms, area behind the plenary area and also a space on the roof.

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.

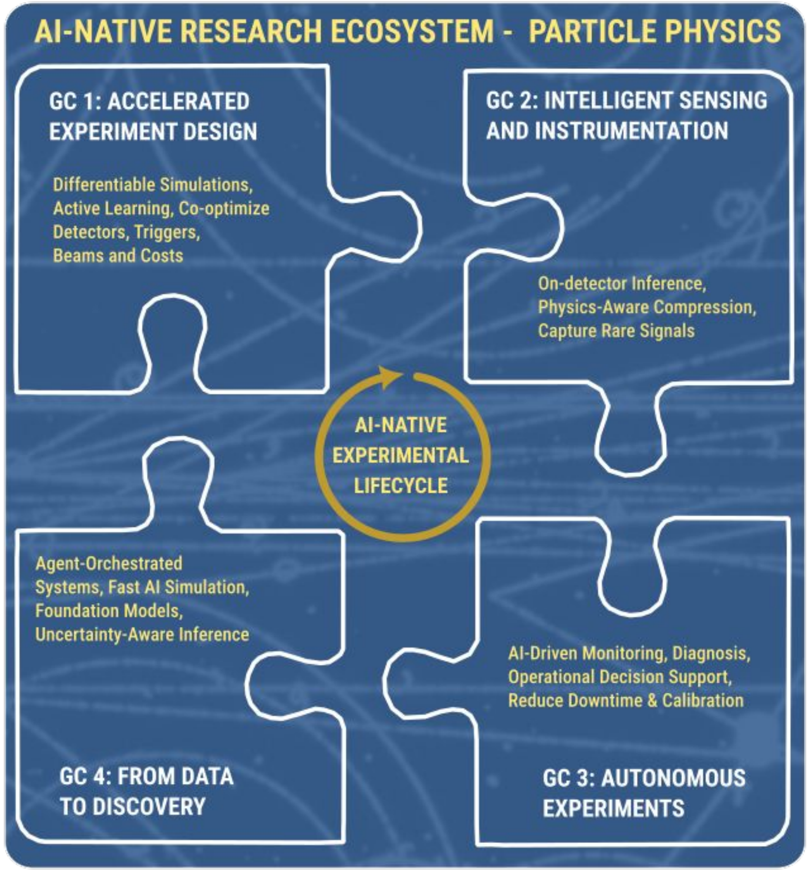

National Scale Collaboration?

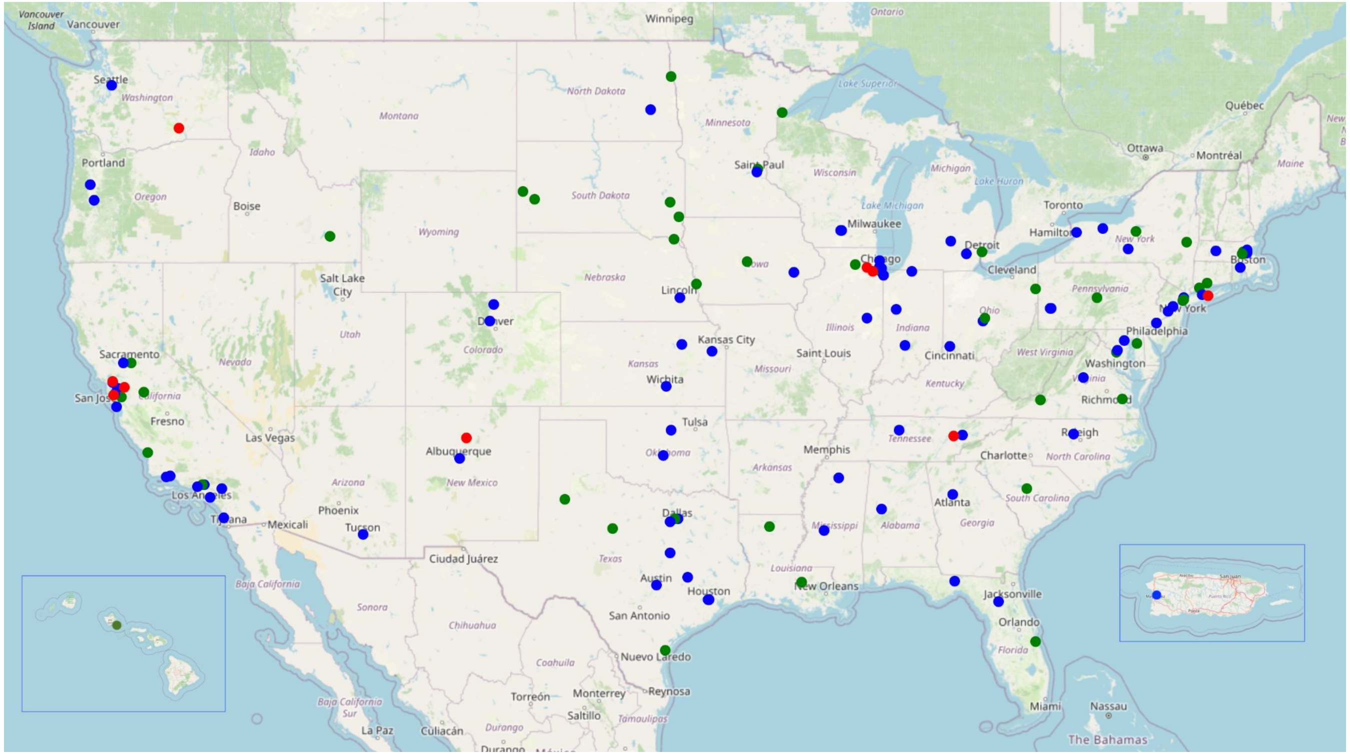

U.S. Institutions (blue dots = universities, red dots = DOE labs) that are currently part of large national-scale funded projects:

- 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.

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.

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.

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.

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.

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.

Discussion

Since the LHC is typically adequately represented in these meetings, we decided to do very short summaries of US-ATLAS and US-CMS together now rather than have (for example) (a) full 10 minute talk(s) for LHC.

US ATLAS

Extensive use of ML/AI across many areas with transformative impact, including particle ID, analysis, reconstruction, simulation, trigger, etc

Jet flavor with transformer architecture

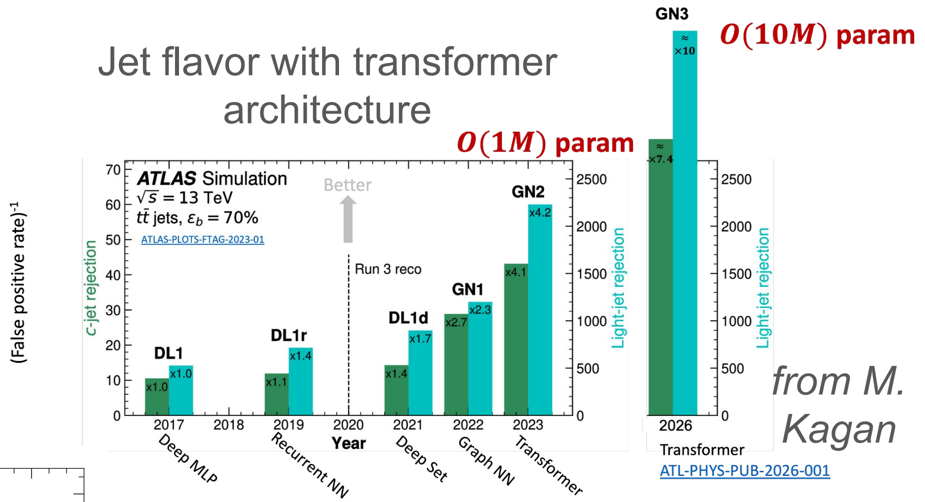

from M. Kagan

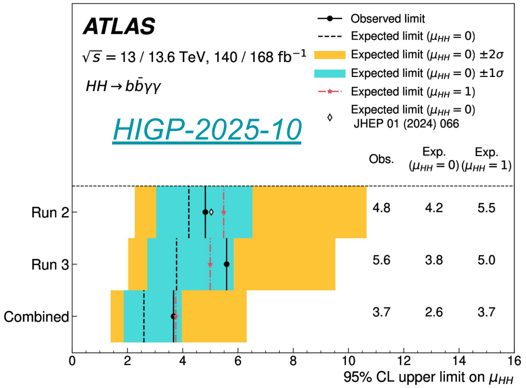

HH → bbγγ Run 2+ Run 3 (2022-2024) comparable to full Run 2 comb
20% sensitivity improvement from new b-tagging capability

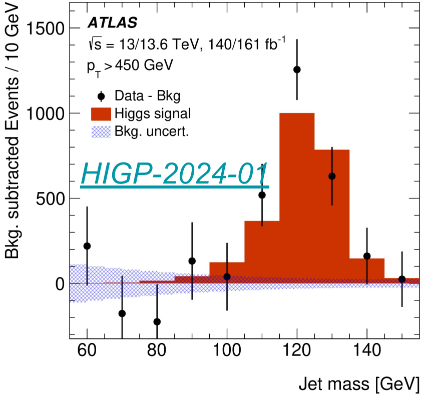

High p_T Higgs measurements
GN2X H → bb tagger
x 7-10 sensitivity improvement compared to Run 2 result

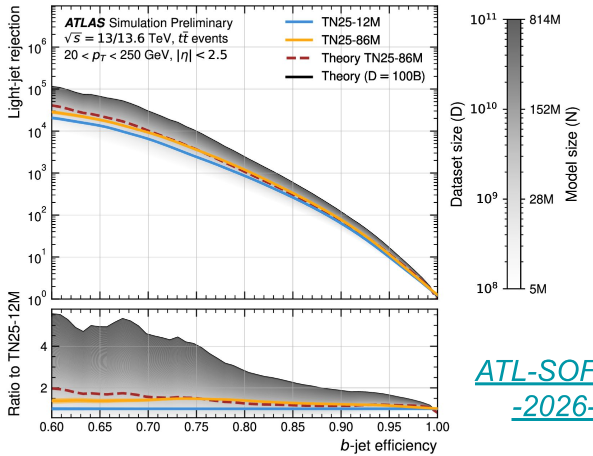

ATL-SOFT-PUB-2026-002

Scaling: improvement by scaling in dataset & model size

US ATLAS

ATLAS requires state-of-the-art AI/ML methods to reach the technical and physics goals for the HL-LHC

- Trigger for new topologies, high precision
- Calibrations, anomaly detection for analysis, NSBI, etc

Agentic methods for detector & computing operations, etc

Investments in infrastructure & tools

- Computational resources & methods (Inference as a Service)
- Frameworks & models

Many projects in the community & interest

- Developments spread across institutes, operations programs, HEP-CCE, IRIS-HEP, A3D3 institutes, Genesis

Need much greater scale to succeed in the future - only possible bringing together the broader community!

Fast simulation with generative methods FastCalo, FATRAS, etc

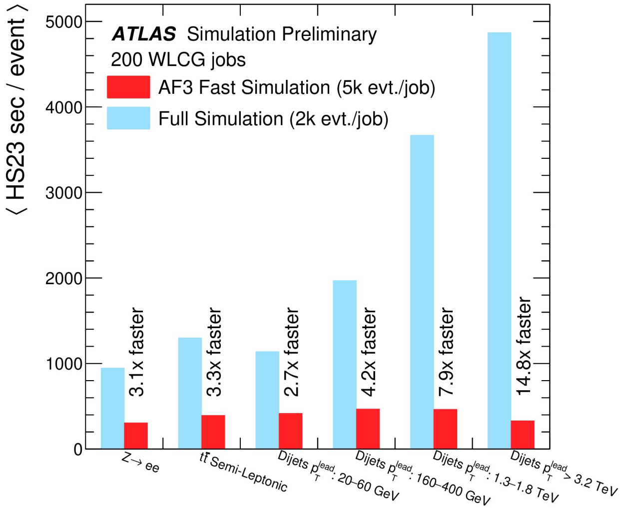

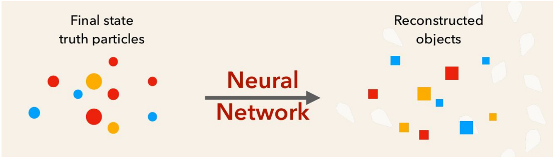

Particle tracking with graph networks

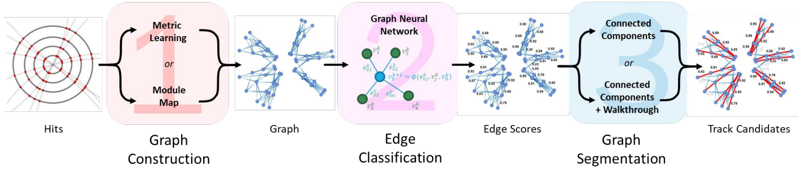

U.S. CMS and AI/ML

- Countless applications in trigger, simulation, and reconstruction (and physics object identification, and analysis)
- Operations: Leveraging LLMs and Agentic AI
  - Computing operations: assistance, documentation, and automation
  - Data Quality Monitoring
  - HEP analysis assistant
- Enabling technology and infrastructure:
  - SuperSONIC enables deployment of ML-inference as-a-service on national cyberinfrastructure
  - ML Training Facility
  - and more…

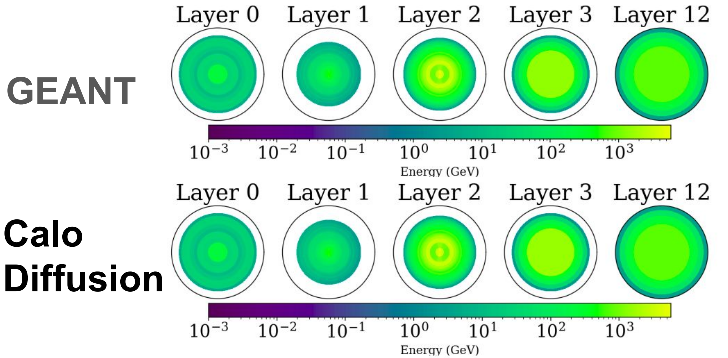

Generative AI for calorimeter simulation

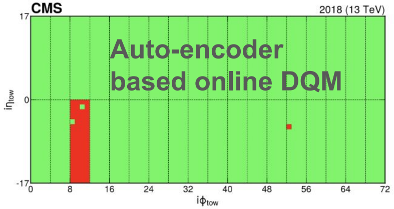

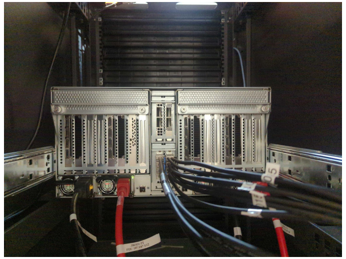

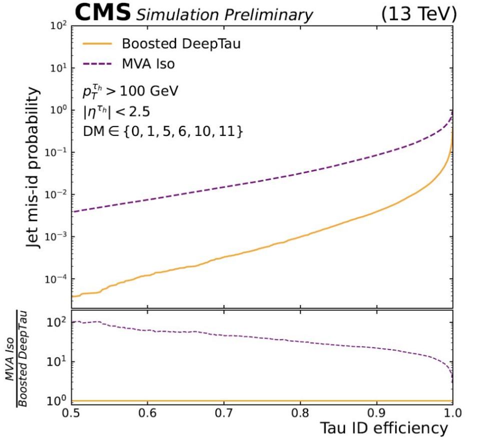

U.S. CMS and the AI-native ecosystem

- Lots of new ideas/projects under discussion – several U.S. CMS members submitted proposals for Genesis RFA
- Topics of interest well-aligned with the four Grand Challenges & Infrastructure
  - AI-driven detector and ASIC design, materials discovery for calorimetry, multi-objective optimization (performance, cost, resilience)...
  - AI in triggers, FPGAs, and on-detector systems, Real-time anomaly detection at 40 MHz, Low-power edge AI…
  - Agentic AI for computing operations and certification, AI-assisted workflow orchestration (incl simulation), Knowledge capture and reuse…
  - Anomaly detection & model-independent searches, Simulation-based inference & EFT analysis, AI cartography of LHC search coverage, Representation learning…
  - Science-aware networks and AI-driven resource optimization, Large-scale distributed computing automation…
- **Strong interest in community-driven approach, facilitation of partnerships, better engagement of universities with DOE labs and industry**

GC1: Accelerated Experimental Design

Use of differentiable/surrogate simulations, and active learning to co-optimize detectors, triggers, beams, and costs, informing design choices for future experimental facilities

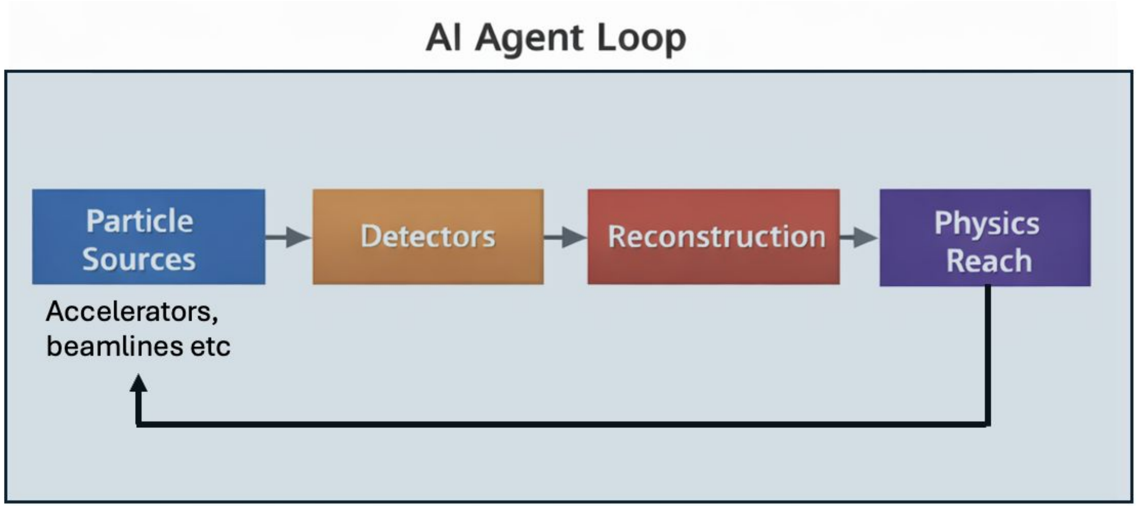

GC2: Intelligent Sensing and Instrumentation

Moving intelligence upstream with on-detector inference, trigger-less/AI-assisted readout, and physics-aware compression to capture rare or unexpected signals without overwhelming bandwidth or storage

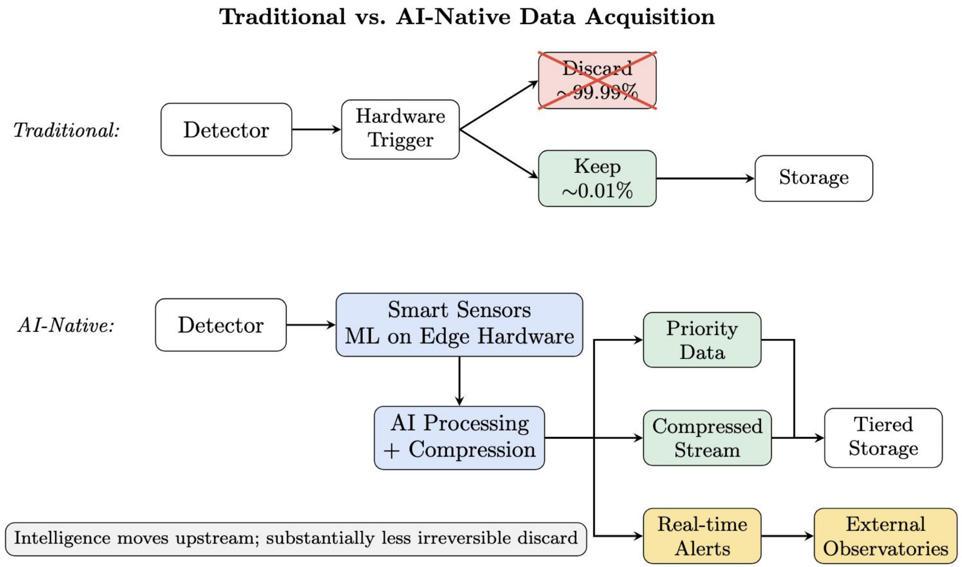

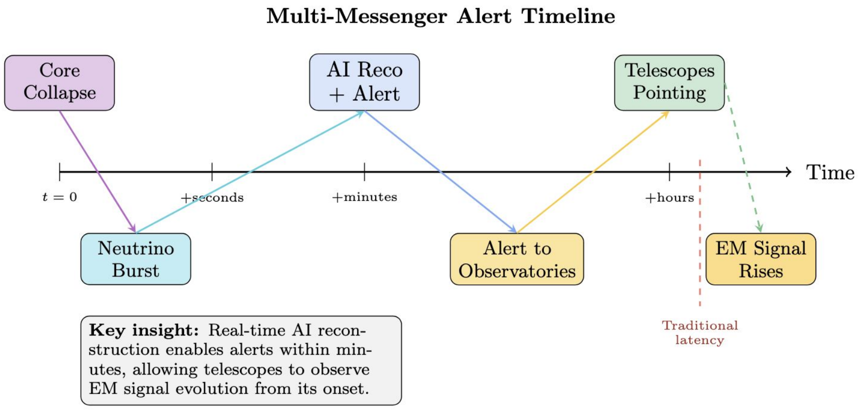

GC3: Autonomous Experiments

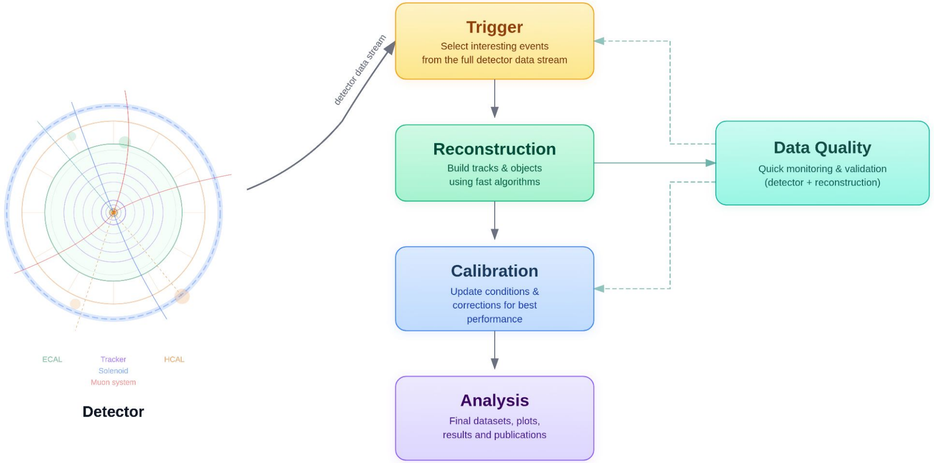

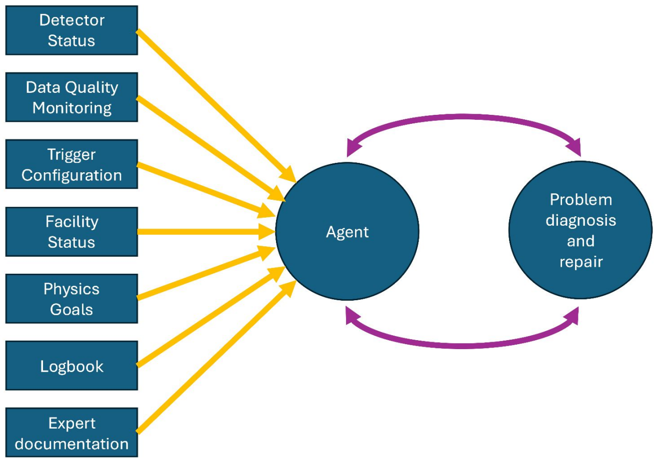

Automating facility operations and calibration with AI-driven monitoring, diagnosis, and operational decision support to reduce downtime, shorten calibration cycles, and preserve institutional knowledge

GC4: From Data to Discovery

Automatic and optimizable data analysis: build agent-orchestrated, goal-directed analysis systems that integrate foundation models, AI-accelerated reconstruction and simulation, and uncertainty-aware inference to dramatically reduce analysis latency, increase scientific productivity, and expand discovery reach

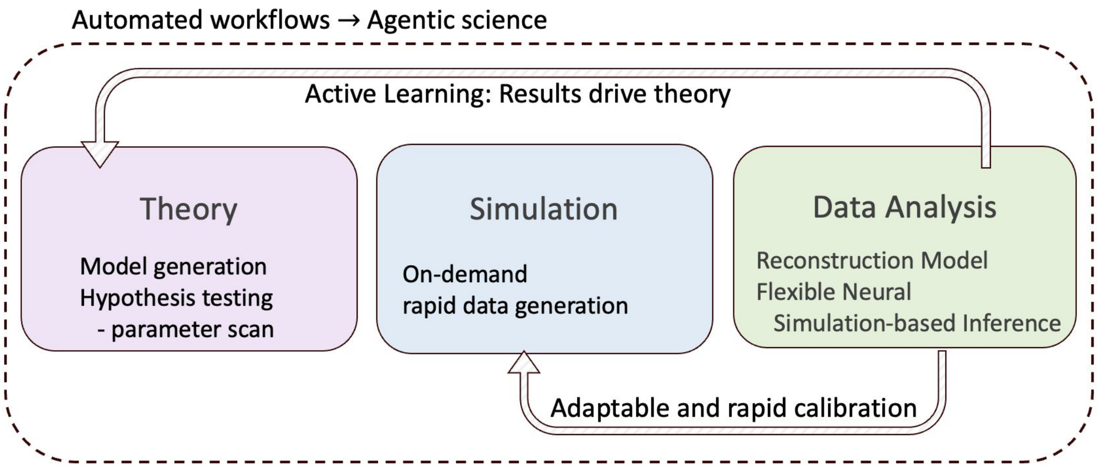

adapted from this talk

Technology, Facilities, Workforce & Community

Development or adaptation of AI technologies for the particle physics setting

Resources to realize the AI vision - includes data access & preparation

Structured pipeline for workforce development

Community comprising a national-scale, multi-institutional collaboration that unites universities, laboratories, and industry partners

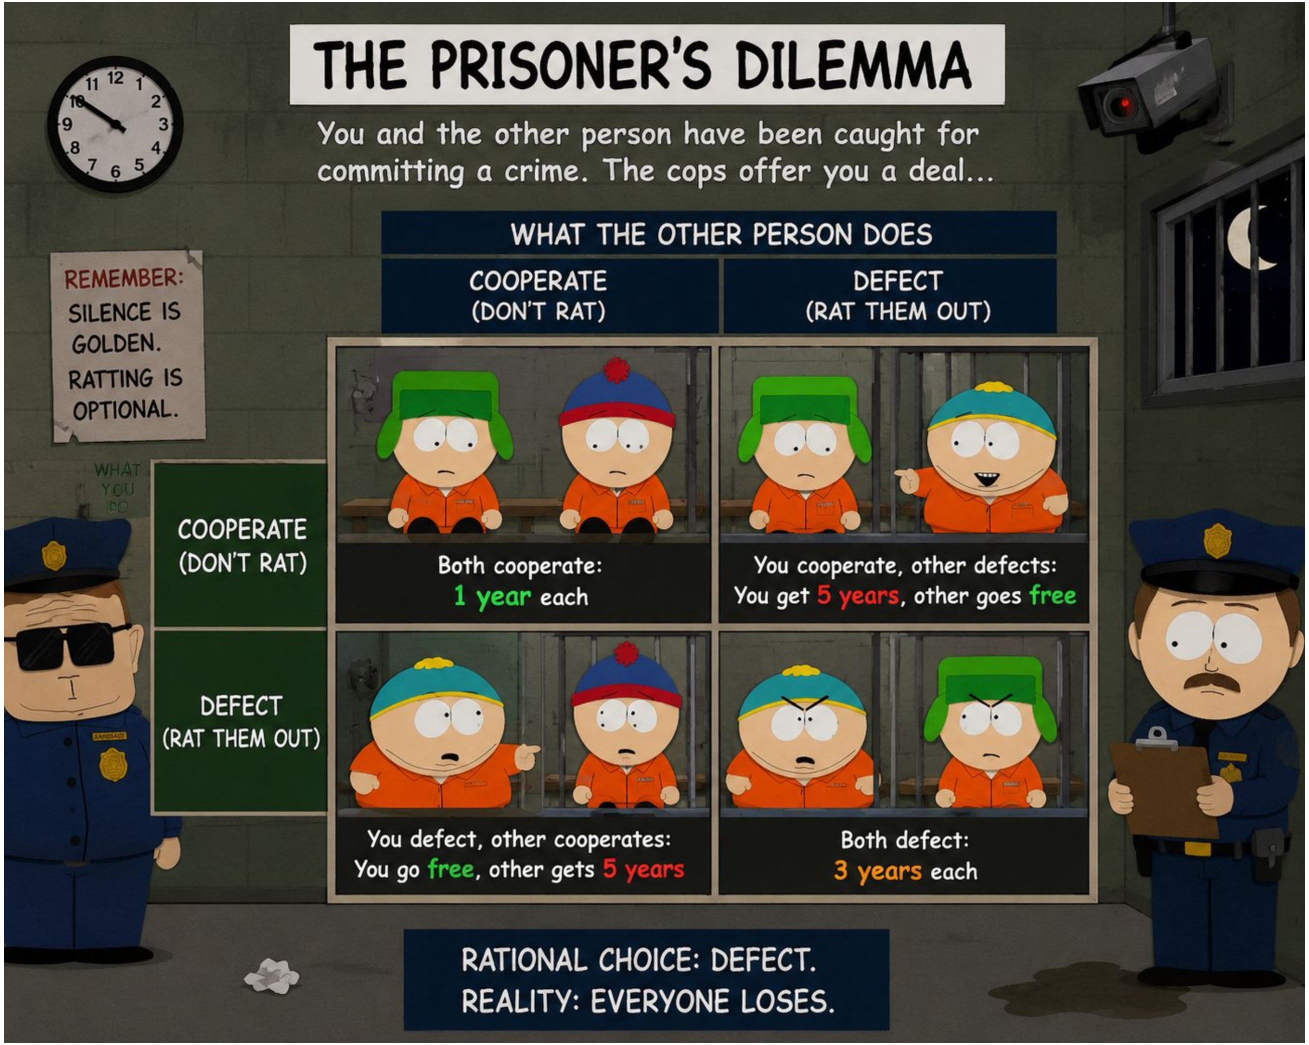
THE PRISONER'S DILEMMA
You and the other person have been caught for committing a crime. The cops offer you a deal...
WHAT THE OTHER PERSON DOES
COOPERATE (DON'T RAT)
DEFECT (RAT THEM OUT)
REMEMBER:
SILENCE IS GOLDEN.
RATTING IS OPTIONAL.
WHAT YOU DO
COOPERATE (DON'T RAT)
DEFECT (RAT THEM OUT)
Both cooperate: 1 year each
You cooperate, other defects: You get 5 years, other goes free
You defect, other cooperates: You go free, other gets 5 years
Both defect: 3 years each
RATIONAL CHOICE: DEFECT.
REALITY: EVERYONE LOSES.