Status of AI/ML at Fermilab

Nhan Tran
for the AI Project Office: Farah Fahim, Burt Holzman, Tia Miceli, Brian Nord, Gabriel Perdue, Tingjun Yang

Fermilab PAC meeting
9 January 2024

Executive summary

for offline reading

Charge: Review the status of the AI/ML activities at the laboratory and of the recommendations made at past meetings: Formulate a strategy to respond to future AI/ML *research* calls, not necessarily just for AI/ML centers.

Framing: AI research is advancing rapidly; one primary area of Fermilab strength is in intelligent sensing and real-time efficient AI

Vision: Accelerate scientific discovery at unprecedented data scales while creating enabling technology for society

Mission: Efficient, robust, autonomous ML codesign

A. Catalyze inclusive, multidisciplinary *Fast ML* community around grand challenges and benchmark tasks
B. Leverage relevant Fermilab core capabilities and strengths to build tools to support the community

Strategy:

A. Identify and grow appropriate sustainable funding streams to support community tools
B. Advance cutting-edge intelligent sensing, real-time AI research
C. Develop industry/academic partnerships to support the core mission

Key performance indicators:

1. Sustainable funding sources for supporting community tools and users on 2 year timescale
2. New and existing partnerships & collaborations resulting in: research output; new projects on AI technology and research; technology transfer; and community growth (users, downloads, etc.)

Outline

- **Framing** - AI in the world, at the DOE
- Dual **Visions** & Fast ML
- **Fast ML Mission**
  - 2 core elements of the strategy
- **Strategy & Key Performance Indicators**

Charge –

We ask the PAC to review the status of the AI/ML activities at the laboratory and of the recommendations made at past meetings: Formulate a strategy to respond to future AI/ML calls, not necessarily just for AI/ML centers.

AI in the world and at the DOE

report

- AI is a big space: industry & academia
- The DOE AI strategy has gone through multiple evolutions
  - Recent excitement on foundation models
  - ChatGPT came out in late 2022
- The leadership class HPC facilities are a key resource (ORNL, ANL, LBNL)
  - most recently – some effort to nucleate around the Trillion Parameter Consortium (TPC)
- **Where does HEP mission and Fermilab fit into this picture?**

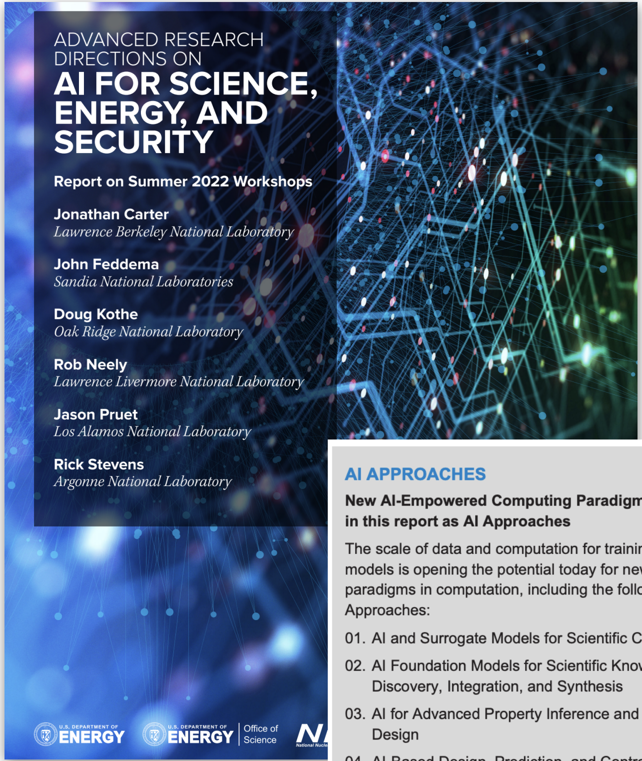

AI APPROACHES

New AI-Empowered Computing Paradigms, known in this report as AI Approaches

The scale of data and computation for training AI models is opening the potential today for new paradigms in computation, including the following AI Approaches:

01. AI and Surrogate Models for Scientific Computing
02. AI Foundation Models for Scientific Knowledge Discovery, Integration, and Synthesis
03. AI for Advanced Property Inference and Inverse Design
04. AI-Based Design, Prediction, and Control of Complex Engineered Systems
05. AI and Robotics for Autonomous Discovery
06. AI for Programming and Software Engineering

Fermilab

Outline

- **Framing** - AI in the world, at the DOE
- Dual **Visions** & Fast ML
- **Fast ML Mission**
- **Strategy & Key Performance Indicators**

Vision

We have consistently focused on:

AI for physics ⇔ physics for AI

- Develop **AI capabilities to accelerate HEP science** <u>and</u> contribute **greater science + industry AI ecosystem**
- Build **diverse, inclusive community; assemble multi-disciplinary collaborations** around cross-cutting HEP AI challenges

Vision

We have consistently focused on:

AI for physics ⇔ physics for AI

This space is massive –

In 2023 PAC talk, we summarized all the exciting activities at Fermilab.

It continues to be our goal to support **all directions of AI research** to advance HEP science, e.g. see recent FNAL AI Jamboree

Vision

We have consistently focused on:

AI for physics ⇔ physics for AI

This space is massive –

In 2023 PAC talk, we summarized all the exciting activities at Fermilab.

It continues to be our goal to support **all directions of AI research** to advance HEP science, e.g. see recent FNAL AI Jamboree

screenshot from last PAC AI talk

AI program in ~15 minutes

- Algorithms for HEP science
  - **Physics-inspired models and data**
    - Graph learning
    - Generative models
    - SBI/likelihood-free inference
    - Accelerating theory
  - **Robust and generalizable learning**
    - Domain adaptation
    - Anomaly detection
    - Semi-/self-supervision
  - **Fast and efficient algorithms**
    - Multi-objective optimization
    - Quantization/sparsity
    - Knowledge distillation
- Operations and controls
  - Real-time accelerator controls
  - Telescope design and operations
  - Quantum machine learning
- Computing hardware and infrastructure
  - Resources for AI practitioners
  - Efficient AI-in-production
- Real-time systems at the edge
  - Hardware-algorithm codesign for HEP and beyond
  - Near-detector, low latency AI
  - On-sensor/detector AI

Fermilab

Vision

We have consistently focused on:

AI for physics ⇔ physics for AI

What is our unique value proposition in the AI space as it pertains to the *PAC charge*?

<u>Accelerate scientific discovery at unprecedented data scales</u>

This will be the focus of this talk.

Vision

We have consistently focused on:

AI for physics ⇔ physics for AI

What is our unique value proposition in the AI space as it pertains to the *PAC charge*?

<u>Accelerate scientific discovery at unprecedented data scales</u>

This will be the focus of this talk.

The strategy:

- is complementary to HPCs (e.g. edge vs cloud)
- leverages unique HEP strengths in cutting-edge sensing technology
- cuts across many scientific domains and industry

Fast ML for Science vision:

Accelerate scientific discovery at unprecedented data scales

> "Scientific discoveries come from groundbreaking ideas and the capability to validate those ideas by testing nature at new scales—finer and more precise temporal and spatial resolution. This is leading to an explosion of data that must be interpreted, and ML is proving a powerful approach. The more efficiently we can test our hypotheses, the faster we can achieve discovery. To fully unleash the power of ML and accelerate discoveries, it is necessary to embed it into our scientific process, into our instruments and detectors."
>
> Applications and Techniques for Fast Machine Learning in Science

Core ML Mission: Efficient, robust, autonomous ML codesign

Outline

- **Framing** - AI in the world, at the DOE
- Dual **Visions** & Fast ML
- **Fast ML Mission**
- **Strategy & Key Performance Indicators**

Fermilab

Mission: Efficient, robust, autonomous ML codesign

- Catalyze **inclusive, multidisciplinary *Fast ML* community** around grand challenges and benchmark tasks
- Leverage strength and scale of national laboratories to **develop critical technologies** that support the community

Grand challenges for HEP, examples

Optimal, continuous readout for DUNE for neutrino physics, multi-messenger astronomy, and other rare measurements

Analyze all 40 MHz of LHC data for the full detector for new physics searches, Higgs measurements, and more

AI-assisted, real-time operation of the [Fermilab accelerator complex]

Fermilab accelerator complex

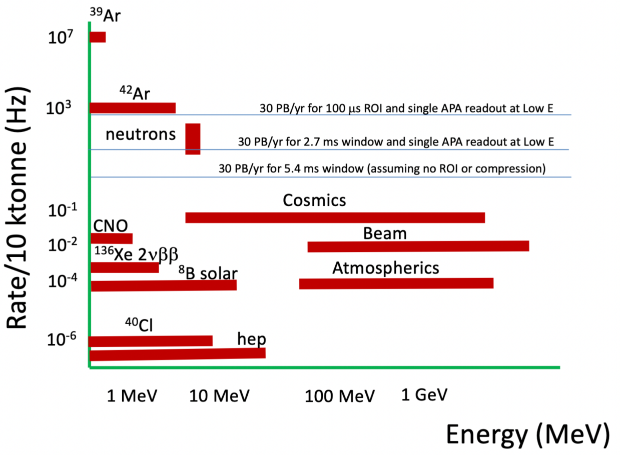

Fermilab

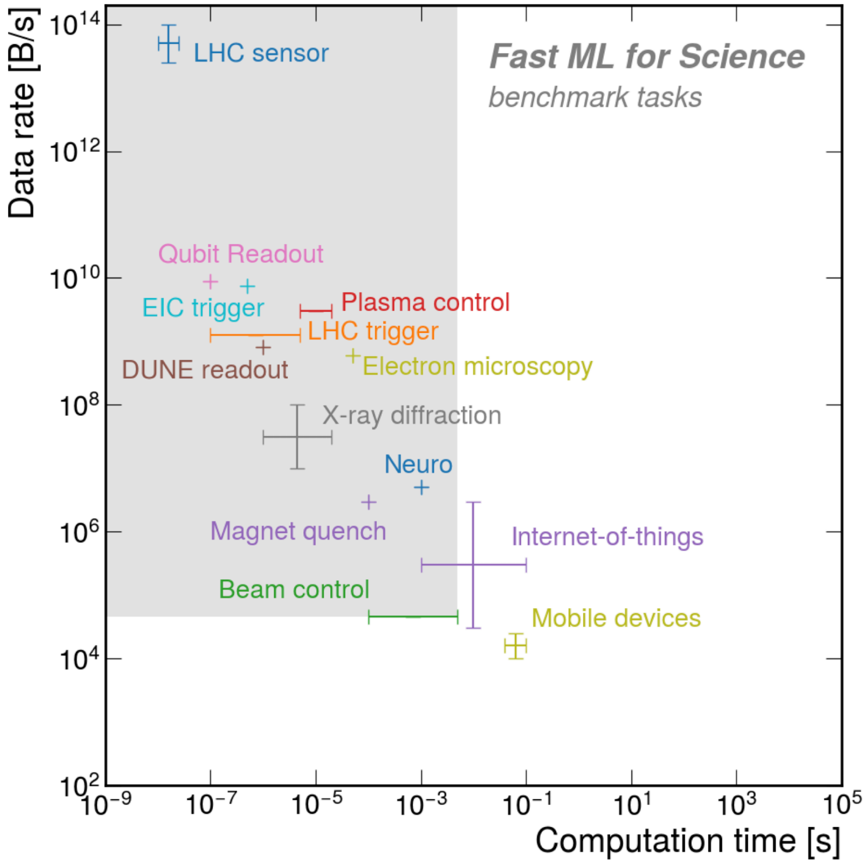

Grand challenges spark imaginations!
Benchmarks bring innovation

Fermilab

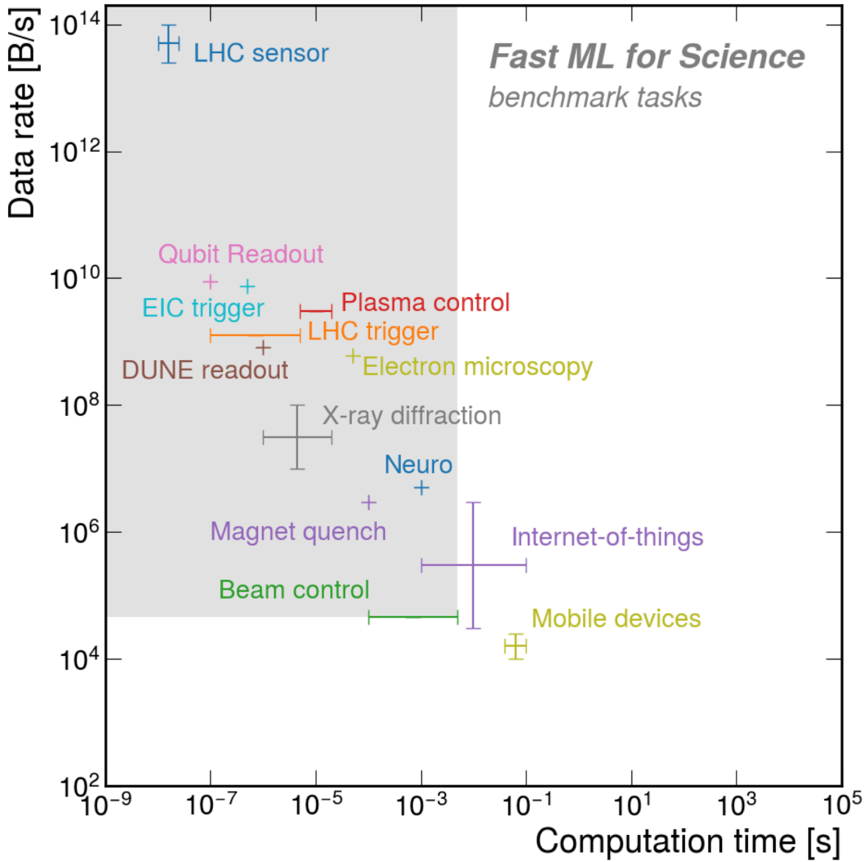

Grand challenges spark imaginations! Benchmarks bring innovation

Benefits to HEP: bring new resources to bear on HEP grand challenges (industry partnerships, computer science & engineering researchers)

HEP-born technology brings **transformative technology** to new material research, fusion energy, neuroscience, or industry applications and so on…

Fermilab

https://fastmachinelearning.org

The Fast ML community

- **Weekly meetings**
- **Annual workshop**
  - 1st edition at FNAL, now on 5th edition in 2024 (bidding for venue now)
  - 1st ICCAD (computer aided design workshop) this year
- **Projects** supported from DOE and NSF coming from the community
  - Example: NSF HDR A3D3 institute
  - DOE funding from HEP, ASCR, NP, SBIR
    - will discuss in more detail later
  - Includes strong local university connections: IIT, Northwestern, Purdue, UC, UIC, UIUC,...
  - Includes international collaborators
- **Connected to wider communities**
  - MLCommons
  - Microelectronics initiatives
  - Industry collaborations and engagements

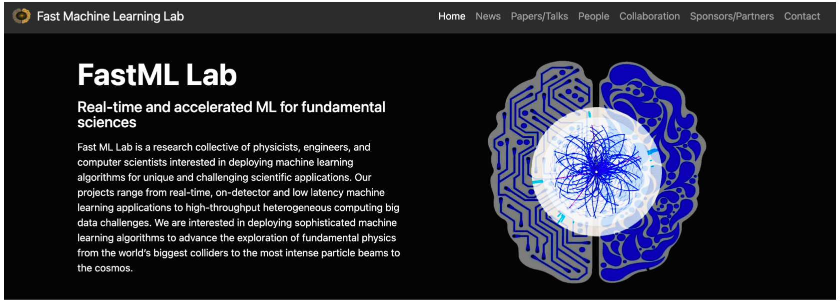

Fermilab

The Fast ML community

Fast ML for Science @ ICCAD, 1st ed. (screenshot)

Time (PDT)	Duration	Presentations		
8:15	15'	Welcome and introduction *Javier Duarte, UCSD*		[Slides]
8:30	30'	Community Vision, Needs, and Progress *Vladimir Loncar, MIT*		[Slides]
9:00	30'	Design Tools Perspective: Catapult + hls4ml for Inference at the Edge *David Burnette, Siemens*		[Slides]
9:30	30'	Designing Hardware for Machine Learning *John Wawrzynek, UC Berkeley*		[Slides]
10:00	30'	Coffee		
10:30	30'	Design Tools Perspective: Mapping ML to the AMD RyzenAI Architecture *Elliott Delaye, AMD*		[Slides]
11:00	30'	Fast ML in the NSF HDR Institute: A3D3 *Shih-Chieh Hsu, UW*		[Slides]
11:30	30'	Real-time ML at the Linac Coherent Light Source *Jana Thayer, SLAC*		[Slides]
12:00	60'	Lunch		
1:20	30'	Robust and Efficient Machine Learning for Mission-Critical Applications *Bhavya Kailkhura, LLNL*		[Slides]
1:50	20'	Quantifying the Efficiency of High-Level Synthesis for Machine Learning Inference *Caroline Johnson (UW), Scott Hauck, Shih-Chieh Hsu, Waiz Khan, Stephany Ayala-Cerna, Geoff Jones, Anatoliy Martynyuk, Matthew Bavier, Oleh Kondratyuk, Trinh Nguyen, Jan Silva, Aidan Short (UW)*	[Paper]	[Slides]

Fast ML for Science, 4th ed. (screenshot)

for offline reading only

Fast Machine Learning at the Large Hadron Collider experiments		*Blackett Laboratory, Lecture Theatre 1, Imperial College London*	09:45 - 10:30
Break		*Blackett Laboratory, Foyer, Imperial College London*	10:30 - 11:00
Fast Machine Learning at European XFEL	*Steve Aplin*	*Blackett Laboratory, Lecture Theatre 1, Imperial College London*	11:00 - 11:45
Bridging AI and biomedicine: Towards AI-driven scientific discoveries	*Maria Brbic*	*Blackett Laboratory, Lecture Theatre 1, Imperial College London*	11:45 - 12:30
Fast Machine Learning for accelerator control	*Karin Rathsman*	*Blackett Laboratory, Lecture Theatre 1, Imperial College London*	09:00 - 09:45
Fast Machine Learning for laser wakefield acceleration	*Matt Streeter*	*Blackett Laboratory, Lecture Theatre 1, Imperial College London*	09:45 - 10:30
Break		*Blackett Laboratory, Foyer, Imperial College London*	10:30 - 11:00
Fast ML for fusion simulation, optimization, and control	*Jonathan Citrin*	*Blackett Laboratory, Lecture Theatre 1, Imperial College London*	11:00 - 11:45
Machine Learning in Exoplanet Characterisation	*Ingo Waldmann et al.*	*Blackett Laboratory, Lecture Theatre 1, Imperial College London*	11:45 - 12:30

Mission: objectives

- Catalyze **inclusive, multidisciplinary *Fast ML* community** around grand challenges and benchmark tasks
- Leverage strength and scale of national laboratories to **develop critical technologies** that support the community

Efficient, robust, autonomous ML codesign

Core Fermilab AI capability, strong synergy with microelectronics initiative

Synergies with different AI research areas across Fermilab, focus area for growth

Efficient, robust, autonomous ML codesign

Efficient {**algorithms**, **tools**, **hardware**, **workflows**, **collaborations**} for ML **codesign**

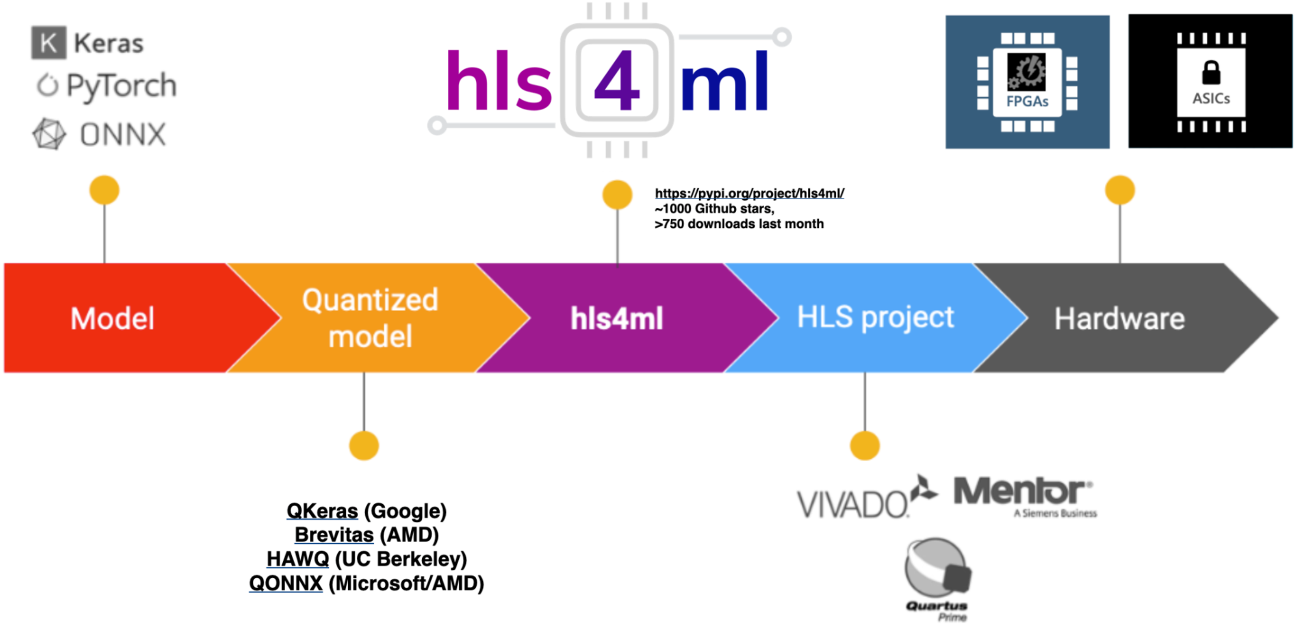

Fermilab

Efficient, robust, autonomous ML codesign

for offline reading only

Efficient {**algorithms**, **tools**, **hardware**, **workflows**, **collaborations**} for ML **codesign**

General techniques: sparsity, quantization
Physics informed techniques: distillation, inductive bias

Efficient and Robust Jet Tagging at the LHC with Knowledge Distillation

Ryan Liu, University of California, Berkeley, Berkeley, CA 94720
Abhijith Gandrakota, Fermi National Accelerator Laboratory, Batavia, IL 60510
Jennifer Ngadiuba, Fermi National Accelerator Laboratory, Batavia, IL 60510
Maria Spiropulu, California Institute of Technology, Pasadena, CA 91125
Jean-Roch Vlimant, California Institute of Technology, Pasadena, CA 91125

AI Engine research collaboration with AMD Architecture group

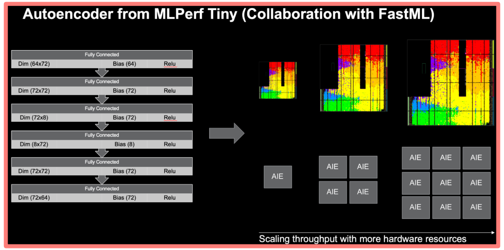

recent ICCAD workshop

Design Tools Perspective: Catapult + HLS4ML for Inference at the Edge

David Burnette
Director of Engineering – Catapult
Siemens EDA

SIEMENS

Short Courses

— **SC1 : Real-time machine learning on FPGAs (hls4ml)**

Course title:	Real-time machine learning on FPGAs (hls4ml)
Course organizer:	Ben Hawks, Fermilab
Date/time/venue:	**Saturday, November 4 – 8:00 am – 6:00 pm – 109 on Level 1**

QONNX: Representing Arbitrary-Precision Quantized Neural Networks

Alessandro Pappalardo, Yaman Umuroglu, Michaela Blott, *AMD Adaptive and Embedded Computing Group (AECG) Labs*, Dublin, Ireland
Jovan Mitrevski[†], Ben Hawks, Nhan Tran, *Fermi National Accelerator Laboratory*, Batavia, IL, USA
Vladimir Loncar[*], *Massachusetts Institute of Technology*, Cambridge, MA, USA
Sioni Summers, *European Organization for Nuclear Research (CERN)*, Geneva, Switzerland
Hendrik Borras, *Heidelberg University*, Heidelberg, Germany
Jules Muhizi, *Harvard University*, Cambridge, MA, USA
Matthew Trahms, Shih-Chieh Hsu, Scott Hauck, *University of Washington*, Seattle, WA, USA
Javier Duarte[†], *University of California San Diego*, La Jolla, CA, USA

with AMD research,
20k downloads last month!

End-to-end codesign of Hessian-aware quantized neural networks for FPGAs and ASICs

JAVIER CAMPOS, JOVAN MITREVSKI, and NHAN TRAN, Fermi National Accelerator Laboratory, USA
ZHEN DONG, AMIR GHOLAMI*, and MICHAEL W. MAHONEY[†], University of California Berkeley, USA
JAVIER DUARTE, University of California San Diego, USA

Efficient, robust, autonom

Efficient {algorithms, tools, hardware, workflov

General techniques: sparsity, quantization
Physics informed techniques: distillation,

Key elements of community for science:

Open-source, not vendor-locked,
community-supported,
user-driven,
accessible and usable

Quantized Neural Networks

Alessandro Pappalardo, Yaman Umuroglu, Michaela Blott
AMD Adaptive and Embedded Computing Group (AECG) Labs
Dublin, Ireland

Jovan Mitrevski[†], Ben Hawks, Nhan Tran
Fermi National Accelerator Laboratory
Batavia, IL, USA

Vladimir Loncar[*]
Massachusetts Institute of Technology
Cambridge, MA, USA

Sioni Summers
European Organization for Nuclear Research (CERN)
Geneva, Switzerland

Hendrik Borras
Heidelberg University
Heidelberg, Germany

Jules Muhizi
Harvard University
Cambridge, MA, USA

Matthew Trahms, Shih-Chieh Hsu, Scott Hauck
University of Washington
Seattle, WA, USA

Javier Duarte[†]
University of California San Diego
La Jolla, CA, USA

with AMD research,
20k downloads last month!

3. **Convert with hls4ml:**

python
```python
import hls4ml

# Configure the converter
config = hls4ml.utils.config_from_keras_model(model, granularity='model')
print(config)

# Convert the model
hls_model = hls4ml.converters.convert_from_keras_model(model,
                                                        hls_config=config,
                                                        output_dir='hls4ml_p
                                                        fpga_part='xcvu9p-fl

# Print the model configuration to check
hls4ml.utils.plot_model(hls_model, show_shapes=True, show_precision=True, t
```

In the `hls4ml.converters.convert_from_keras_model` function:

- `hls_config`: is the configuration for the conversion. We generated it using `config_from_keras_model` for simplicity, but you can customize this as needed.
- `output_dir`: is the directory where the HLS project will be created.
- `fpga_part`: specifies the FPGA part. Adjust this based on your target FPGA.

4. **Build the HLS Project:**

Once the model is converted, you can compile it into HLS:

python Copy code
```python
hls_model.compile()
```

5. **Run the HLS Simulation (Optional):**

Before synthesizing for FPGA, you can run a C-simulation to check if the model works correctly in HLS:

python Copy code
```python
hls_model.build(csim=True)
```

After this, you'll have an HLS project in the specified `output_dir` that you can use with FPGA development tools to generate bitstreams for FPGA deployment.

for offline reading only

kshop

ve:

IEMENS

on Level 1

USA

keley,

ilab

Jindariani, Ngadiuba, Pedro, Tran, Comput Softw Big Sci (2019) 3:13
Kljinsma, Pedro, Tran, Mach. Learn.: Sci. Technol. 2 (2021) 035005
Kljinsma, Pedro, Tran, IEEE/ACM H2RC 2020
Wang,Yang,Flechas,Hawks,Holzman,Knoepfel,Pedro,Tran, arXiv:2009.04509
Cai,Herner,Yang,Wang,Flechas,Holzman,Pedro,Tran, arXiv:2301.04633

Efficient, robust, autonomous ML codesign

for offline reading only

Efficient {**algorithms**, **tools**, **hardware**, **workflows**, **collaborations**} for ML **codesign**

Another example: system codesign for heterogeneous grid computing to accelerate ML workflows

- To alleviate future HEP computing will be bottlenecks - **enable more powerful algorithms** on optimal hardware
- **Coprocessors** (GPUs, FPGAs, ASICs, …) naturally accelerate ML workloads **by orders of magnitude**
- Leverage industry hardware and tools - provide **coprocessors as-a-service**
- **SONIC**: Services for Optimized Network Inference on Coprocessors

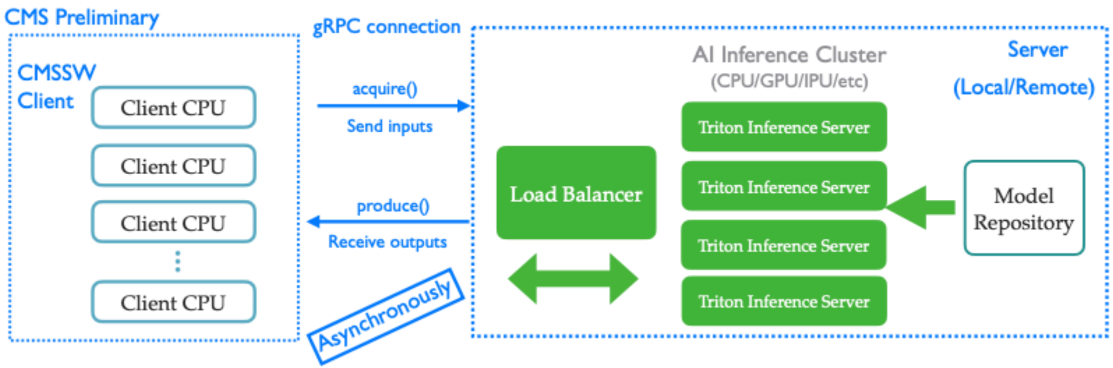

Portable Acceleration of CMS Production Workflow with Coprocessors as a Service

The CMS Collaboration

MLG-23-001, first CMS ML Group paper!

for offline reading only

Efficient, robust, autonomous ML codesign

Example images of simulated galaxy morphologies with different levels of telescope noise.

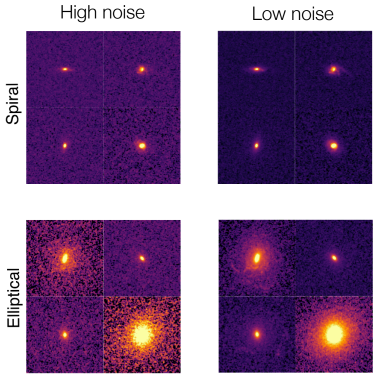

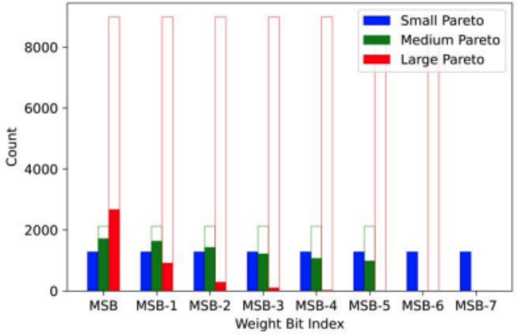

Bit level fault sensitivity analysis

Linac unsupervised fault clustering

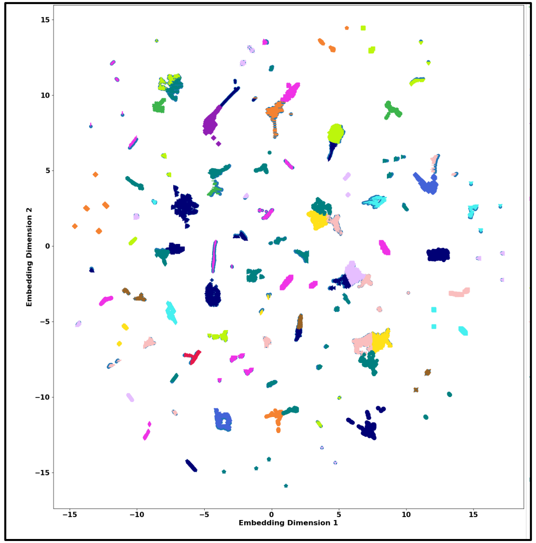

Anomaly detection algorithm running online test bench on CMS Run 3 data!

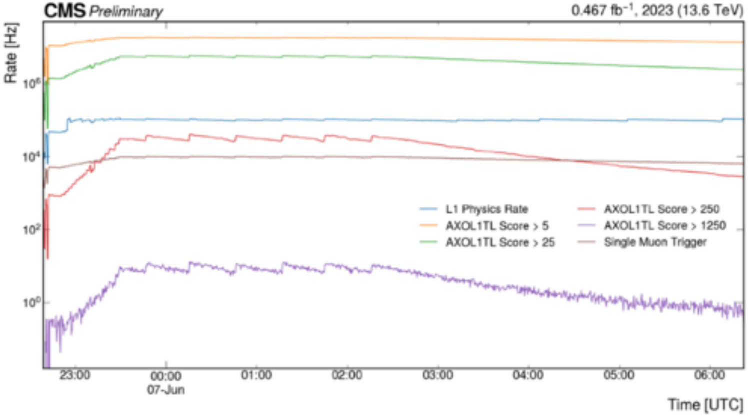

Fermilab

Impact: recent *examples*

Leverage core capabilities to deploy **ML <u>at scale</u>** - algorithms + facilities, tools, software, multidisciplinary teams

e.g. large scale user facilities & advanced instrumentation; advanced computer science, visualization, & data; microelectronics

Impact: recent *examples*

Leverage core capabilities to deploy **ML <u>at scale</u>** - algorithms + facilities, tools, software, multidisciplinary teams

e.g. large scale user facilities & advanced instrumentation; advanced computer science, visualization, & data; microelectronics

- First ever L1 trigger anomaly detection algorithm deployed for LHC CMS Run 3
  - Growth from community benchmarks and collaborations built from community efforts, investment in hls4ml (FastML, AMD, Siemens)
- CMS MLG-23-001 demonstration of accelerated ML workflows with SONIC; working with NVidia, Graphcore, computing operations experts
- First edge AI deployed in Fermilab accelerator complex; working with Intel/NU

https://cds.cern.ch/record/2876546

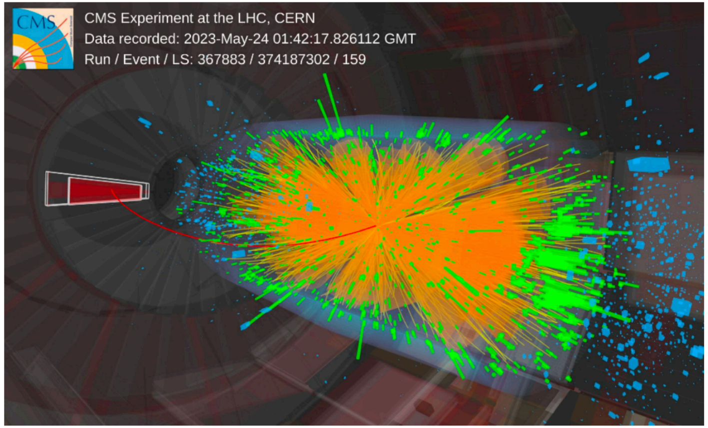

Rameika, [HEPAP Aug23]

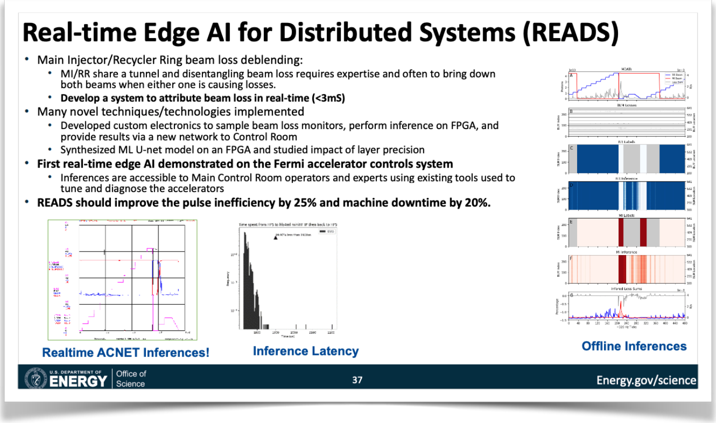

Fermilab

Impact: recent *examples*

for offline reading only

Leverage core capabilities to deploy **ML <u>at scale</u>** - algorithms + facilities, tools, software, multidisciplinary teams

e.g. large scale user facilities & advanced instrumentation; advanced computer science, visualization, & data; microelectronics

Low latency optical-based mode tracking with machine learning deployed on FPGAs on a tokamak

Y. Wei,[1, a)] R. F. Forelli,[2, 3, b)] C. Hansen,[1] J. P. Levesque,[1] N. Tran,[2,4] J. C. Agar,[5] G. Di Guglielmo,[6, 4] M. E. Mauel,[1] and G. A. Navratil[1]
[1)] *Department of Applied Physics and Applied Mathematics, Columbia University*
[2)] *Real-time Processing Systems Division, Fermi National Accelerator Laboratory*
[3)] *Department of Electrical and Computer Engineering, Lehigh University*
[4)] *Department of Electrical and Computer Engineering, Northwestern University*
[5)] *Department of Mechanical Engineering and Mechanics, Drexel University*
[6)] *Microelectronics Division, Fermi National Accelerator Laboratory*

arXiv: 2312.00128

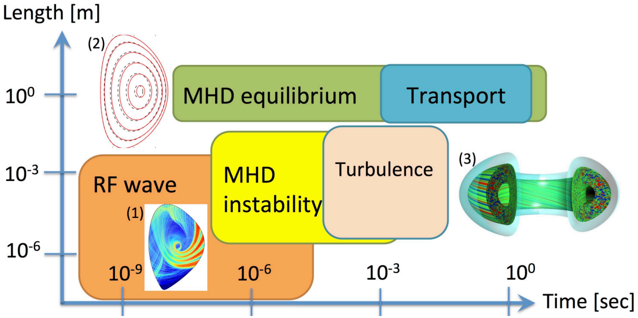

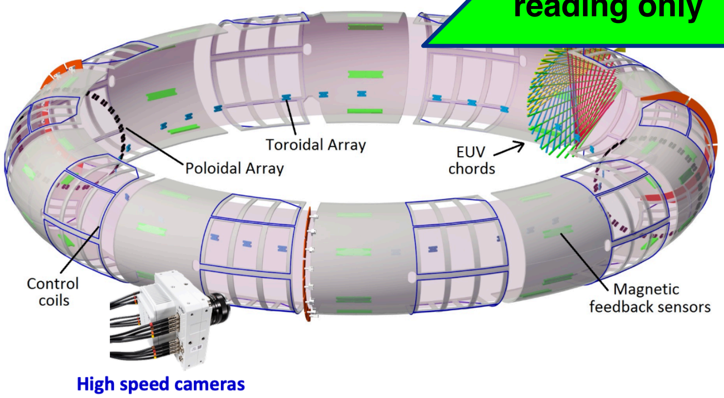

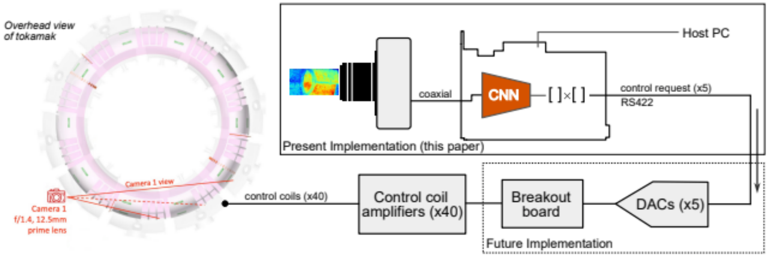

120 kfps throughput, 17.6 μs latency

Enabling new capabilities for fusion experiments!

Outline

- **Framing** - AI in the world, at the DOE
- Dual **Missions** & Fast ML
- **Fast ML Vision**
- **Strategy & Key Performance Indicators**

Fast ML ecosystem

Robust ML
(anomaly detection, domain adaptation, uncertainty quantification,..)

Efficient ML hardware codesign [hls4ml]

Adaptive learning

Fast ML for Science:
autonomous discovery at unprecedented data scales
Efficient, robust, autonomous ML system codesign

Fast ML ecosystem

Robust ML
(anomaly detection, domain adaptation, uncertainty quantification,..)

Efficient ML hardware codesign [hls4ml]

Adaptive learning

Fast ML for Science:
autonomous discovery at unprecedented data scales
Efficient, robust, autonomous ML system codesign

Scientific and industrial edge applications
HEP, NP, BES, FES, QIS, MMA, Neuro, …
internet-of-things, industry 4.0, autonomous vehicles

Fast ML ecosystem

Foundation Models,
Digital Twins

Core ML research

Microelectronics

HPC (supercomputing)

Robust ML
(anomaly detection, domain adaptation, uncertainty quantification,..)

Efficient ML hardware codesign [hls4ml]

Adaptive learning

Fast ML for Science:
autonomous discovery at unprecedented data scales
Efficient, robust, autonomous ML system codesign

Scientific and industrial edge applications
HEP, NP, BES, FES, QIS, MMA, Neuro, …
internet-of-things, industry 4.0, autonomous vehicles

Fast ML ecosystem

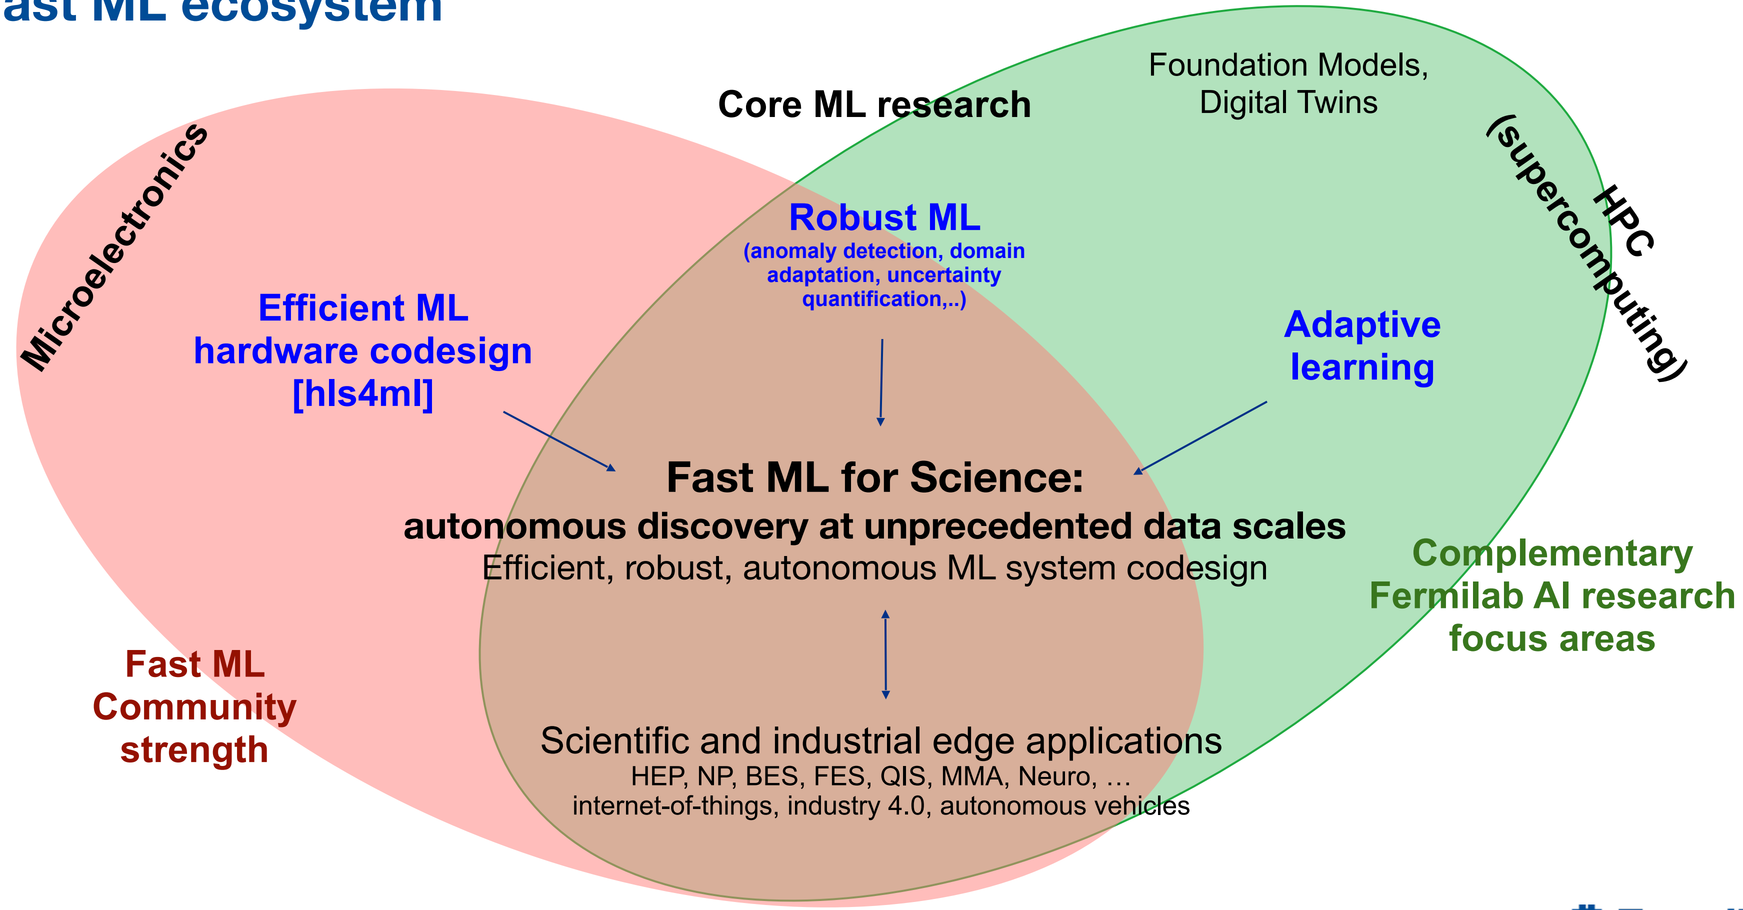

Project support and funding strategy

Support from DOE and community focused around research topics

- DOE HEP (Lab and University awards), ASCR, NP
- Additional sources: JTFI (UChicago/ANL), DPI (UIllinois system), LDRD
- NSF include HDR institute (A3D3)

Strategic directions for support:

A. Identify and develop sustainable funding streams to support broadening community tools and techniques
B. Advance cutting-edge intelligent sensing, real-time AI research and hardware codesign
C. Develop strategic industry/academic partnerships to support the core mission

Strategic growth

Fast ML for Science:
autonomous discovery at unprecedented data scales

Efficient ML codesign	**Robust ML**	**Adaptive learning**

Strategic direction (A)

← Strategic direction (B,C) →

Fermilab

Strategic growth

Fast ML for Science:
autonomous discovery at unprecedented data scales

Efficient ML codesign | **Robust ML** | **Adaptive learning**

Strategic direction (A)

Strategic direction (B,C)

Grand challenges in HEP:

- Optimal, continuous readout for DUNE for neutrino physics, multi-messenger astronomy, and other rare measurements
- Analyze all 40 MHz of LHC data for the full detector for new physics searches, Higgs measurements, and more
- AI-assisted, real-time operation of the Fermilab accelerator complex

And more in HEP and beyond!

Executive summary

for offline reading

Charge: Review the status of the AI/ML activities at the laboratory and of the recommendations made at past meetings: Formulate a strategy to respond to future AI/ML *research* calls, not necessarily just for AI/ML centers.

Framing: AI research is advancing rapidly; one primary area of Fermilab strength is in intelligent sensing and real-time efficient AI

Vision: Accelerate scientific discovery at unprecedented data scales while creating enabling technology for society

Mission: Efficient, robust, autonomous ML codesign

A. Catalyze inclusive, multidisciplinary *Fast ML* community around grand challenges and benchmark tasks
B. Leverage relevant Fermilab core capabilities and strengths to build tools to support the community

Strategy:

A. Identify and grow appropriate sustainable funding streams to support community tools
B. Advance cutting-edge intelligent sensing, real-time AI research
C. Develop industry/academic partnerships to support the core mission

Key performance indicators:

1. Sustainable funding sources for supporting community tools and users on 2 year timescale
2. New and existing partnerships & collaborations resulting in: research output; new projects on AI technology and research; technology transfer; and community growth (users, downloads, etc.)

Additional material

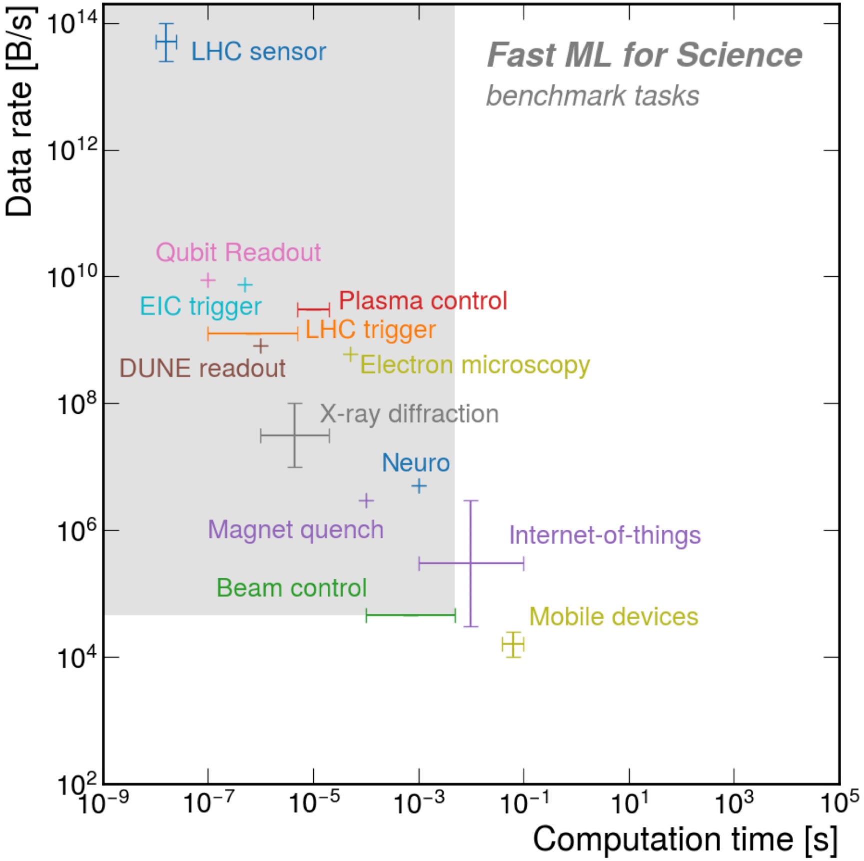

Benchmarks bring innovation, Grand challenges spark imaginations!

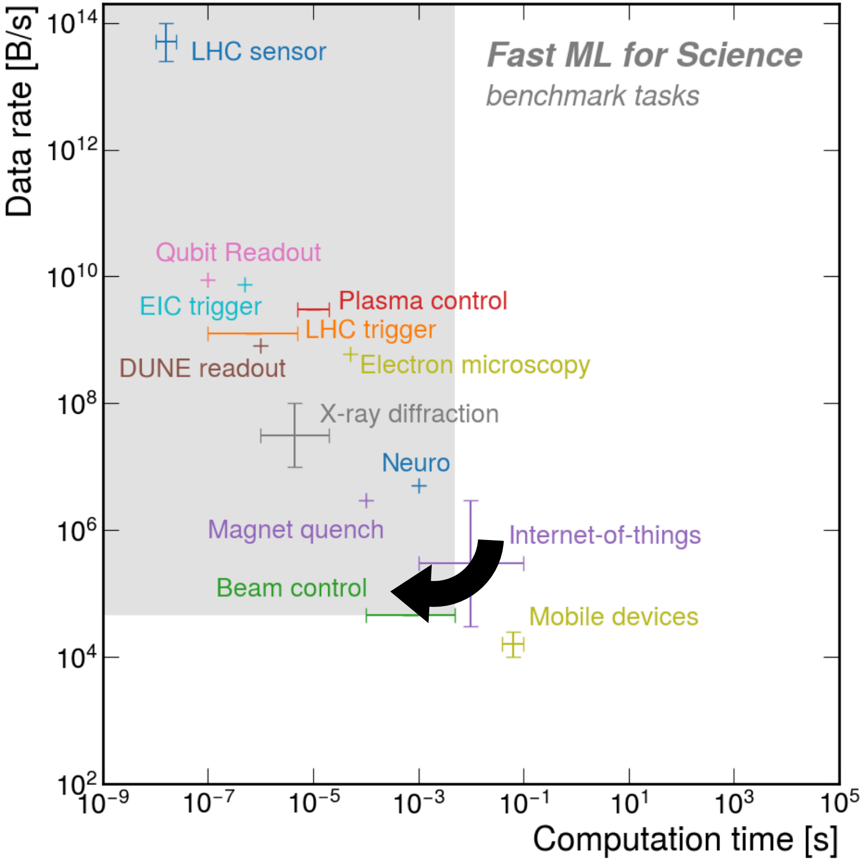

Benchmarks bring innovation, Grand challenges spark imaginations!

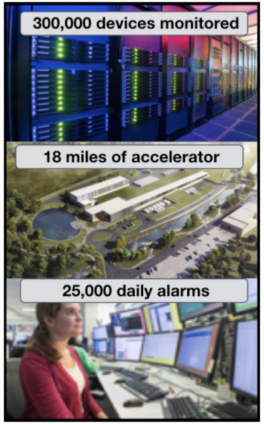

Fermilab accelerator complex

Particle accelerator controls

[Talk by J. Mitrevski]()

Fermilab

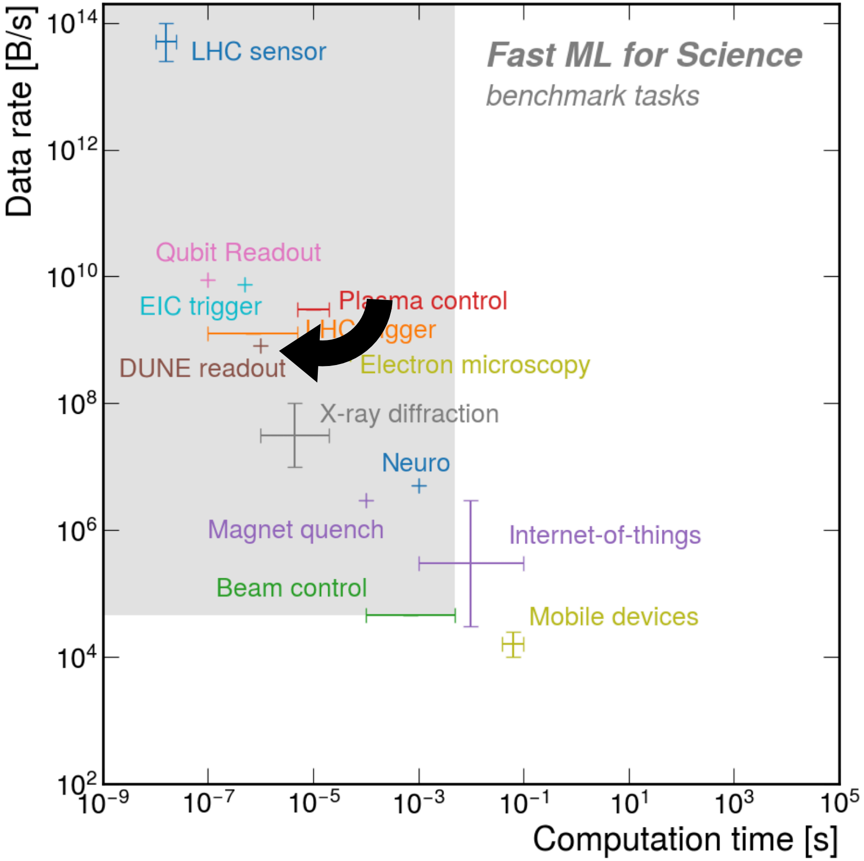

Benchmarks bring innovation, Grand challenges spark imaginations!

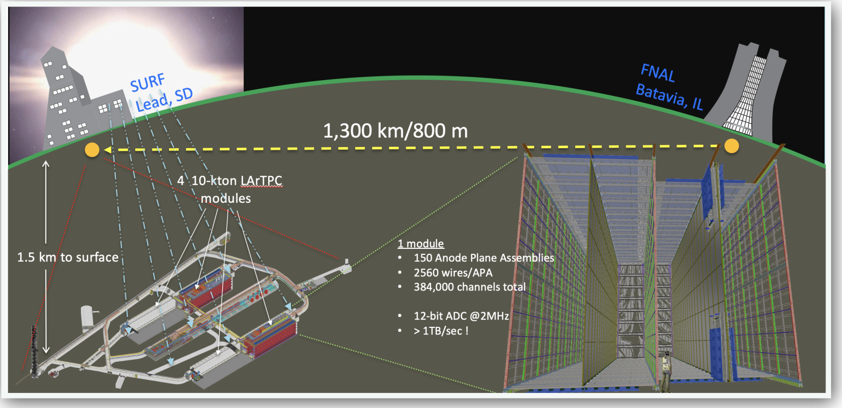

Supernova detection and multi-messenger astronomy
[Talk by M. Kahn]

Fermilab

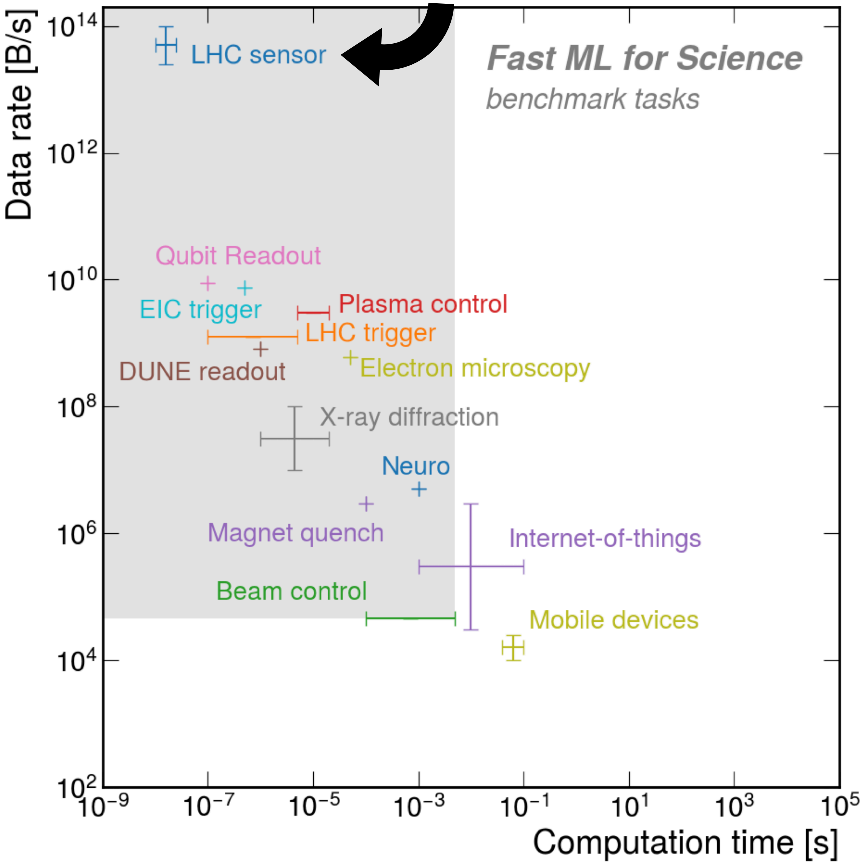

Benchmarks bring innovation, Grand challenges spark imaginations!

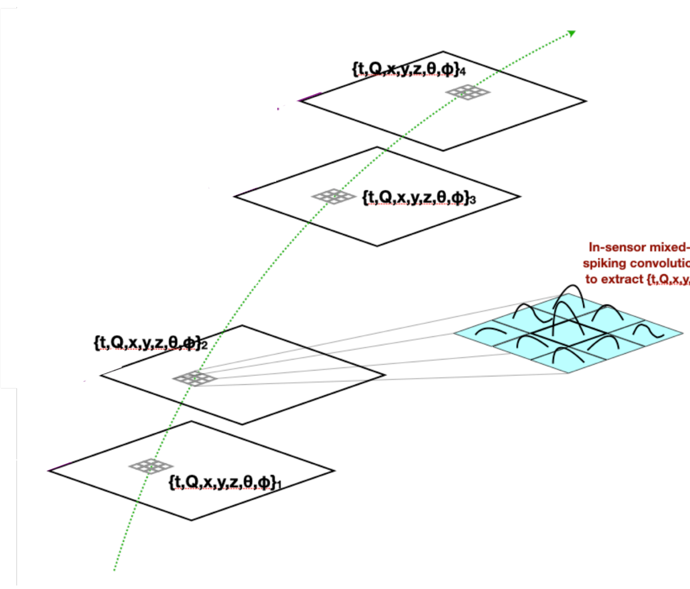

Full 40 MHz readout with smart pixel detectors

[Talk by G. Di Guglielmo]

Fermilab

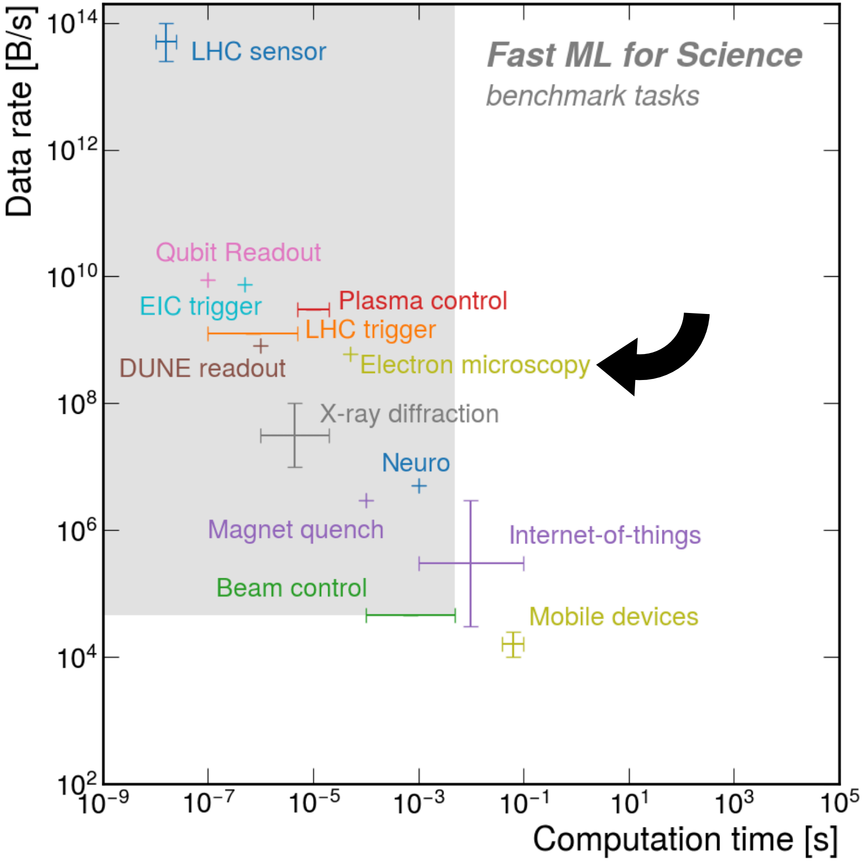

Benchmarks bring innovation, Grand challenges spark imaginations!

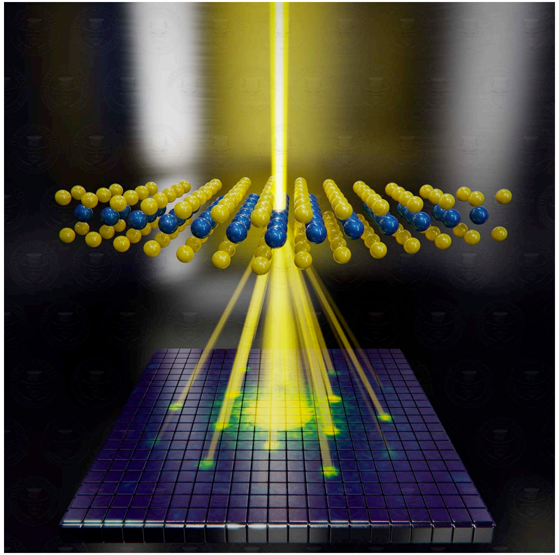

New materials for quantum and energy

[Talk by J. Agar]

Benchmarks bring innovation, Grand challenges spark imaginations!

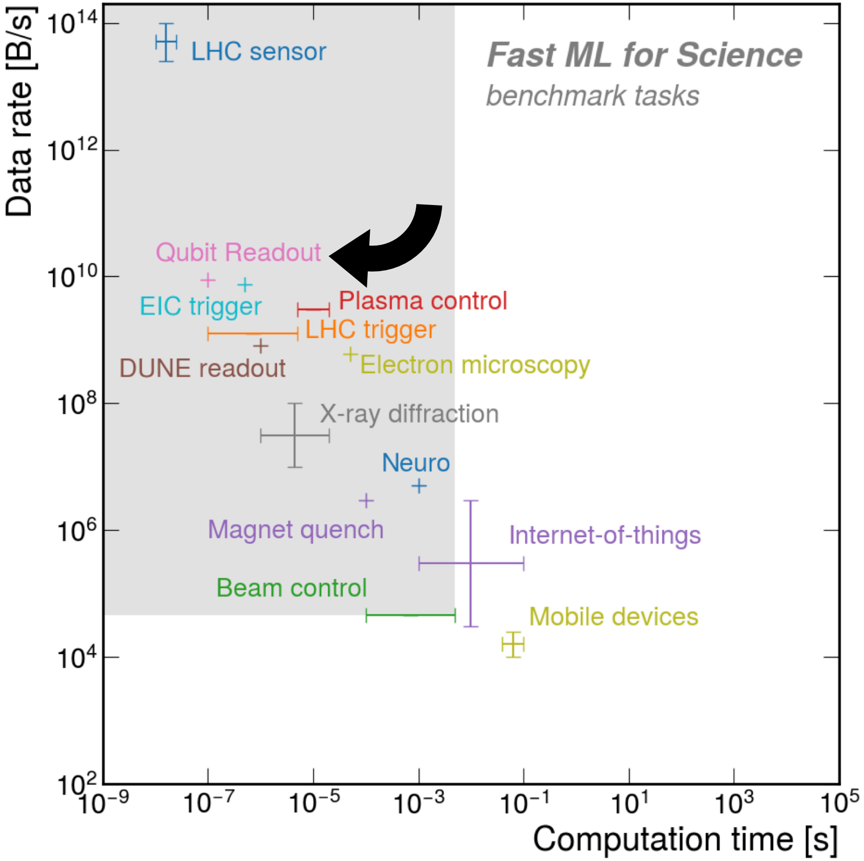

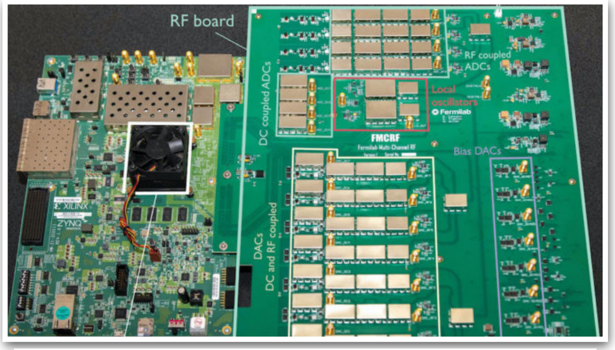

QICK

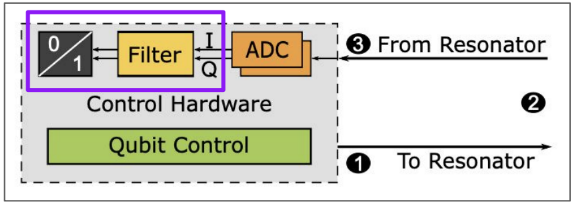

Qubit readout and control
Talk by J. Campos

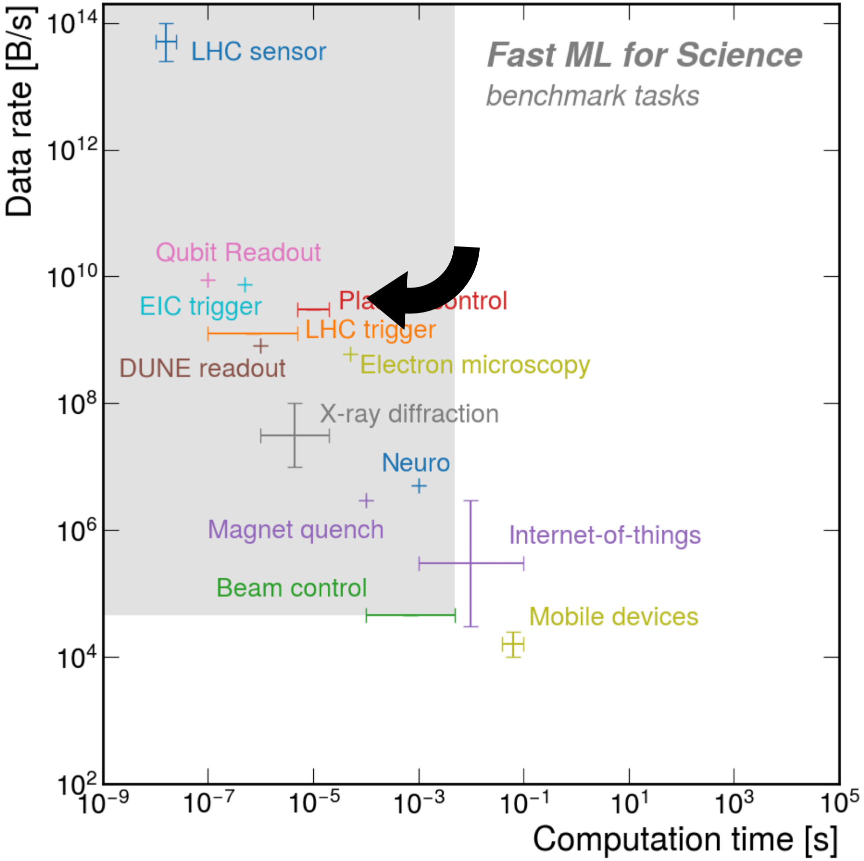

Benchmarks bring innovation, Grand challenges spark imaginations!

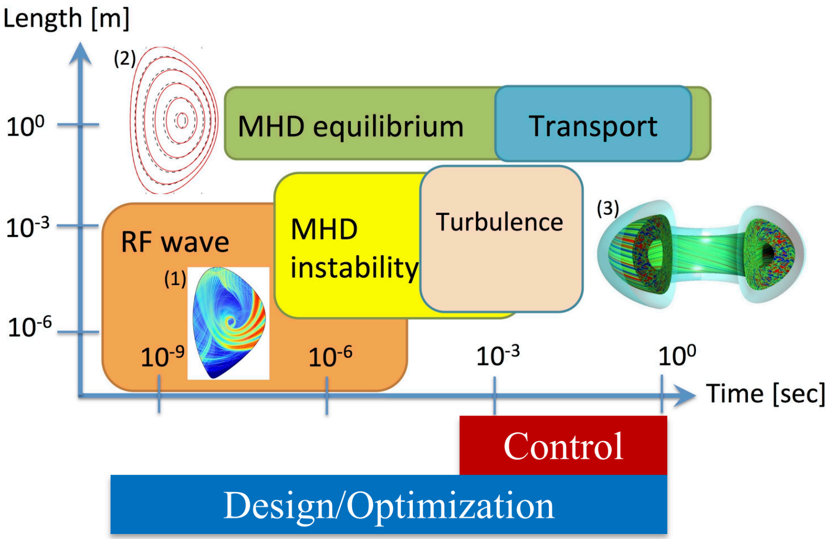

Magnetohydrodynamics Instabilities

Talk by R. Forelli

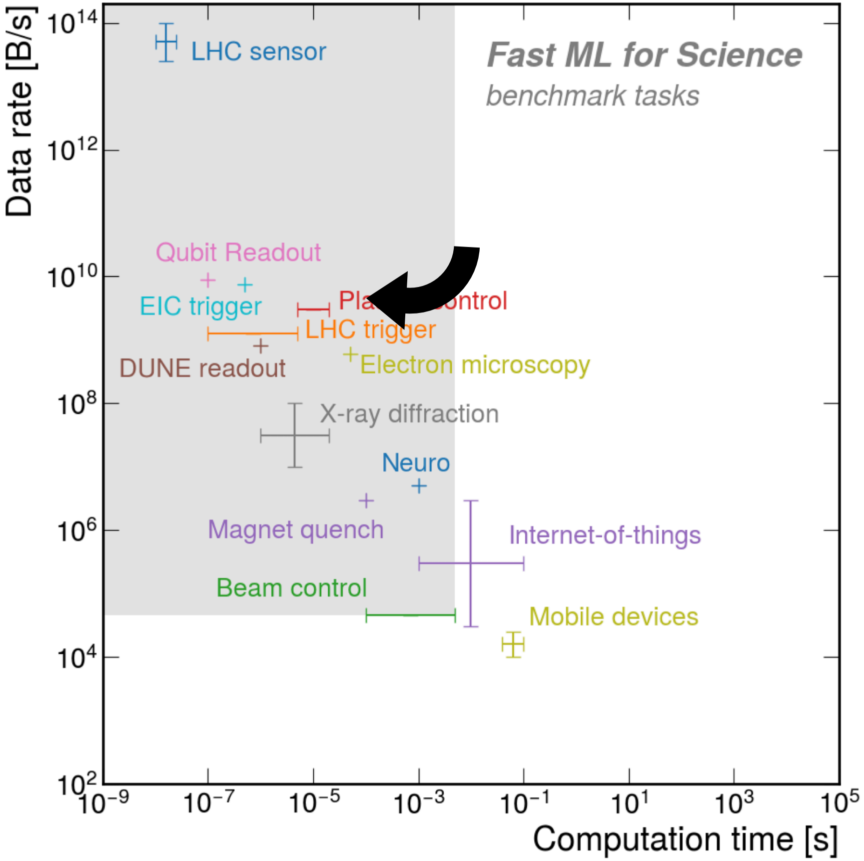

Benchmarks bring innovation, Grand challenges spark imaginations!

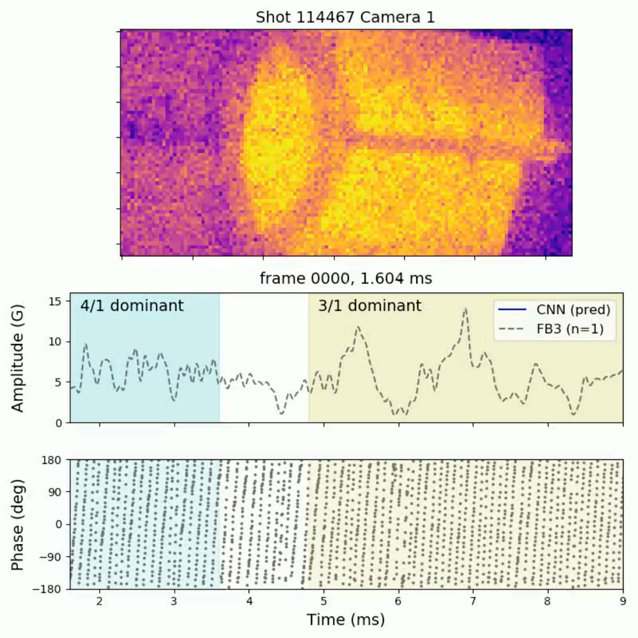

Magnetohydrodynamics Instabilities

[Talk by R. Forelli]

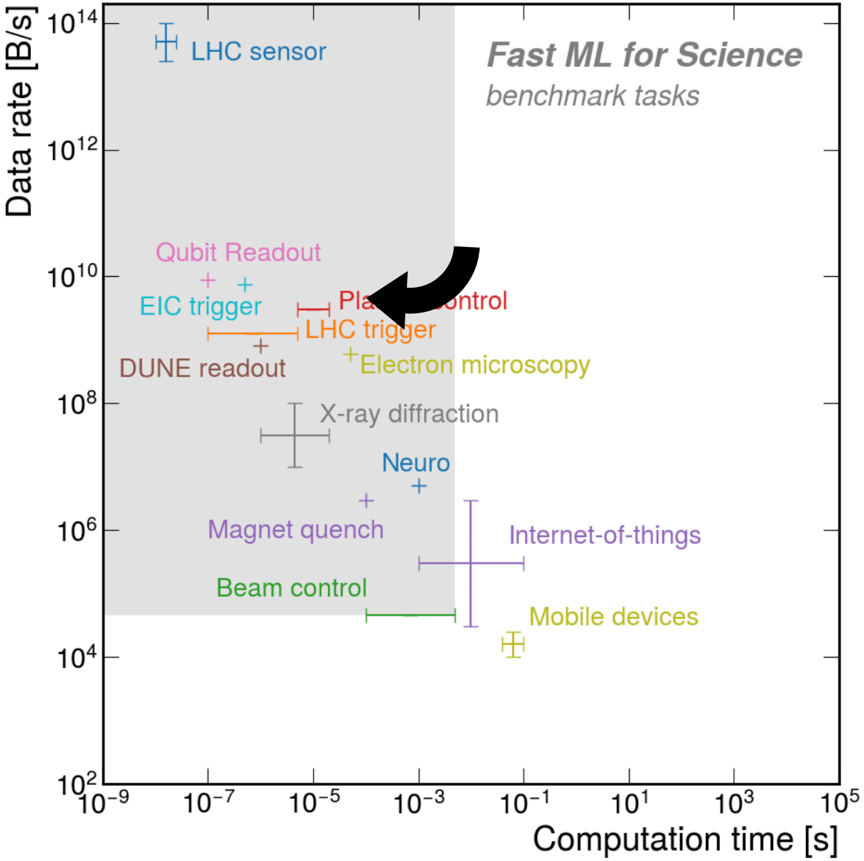

Benchmarks bring innovation, Grand challenges spark imaginations!

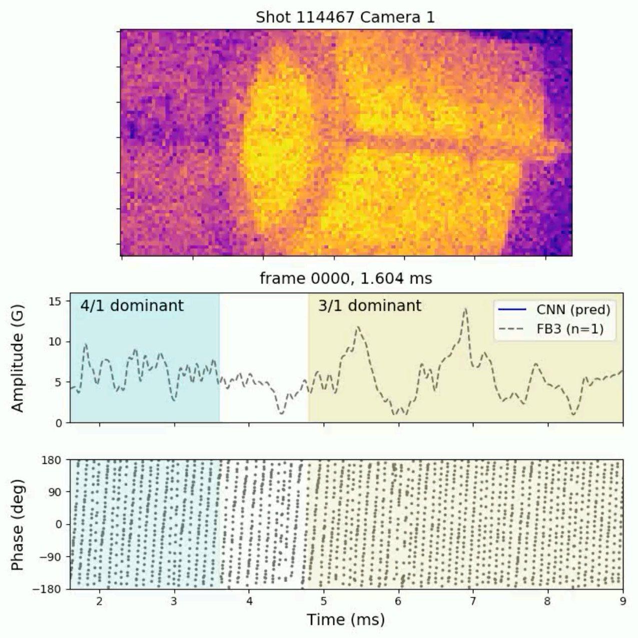

Magnetohydrodynamics Instabilities

[Talk by R. Forelli]

Benchmarks bring innovation, Grand challenges spark imaginations!

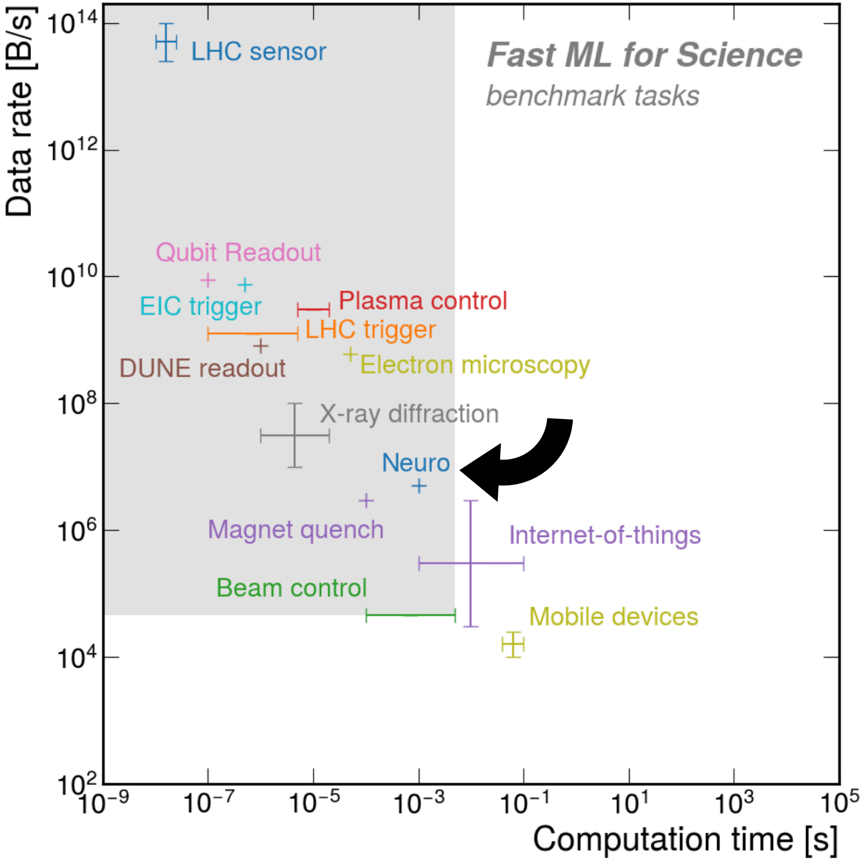

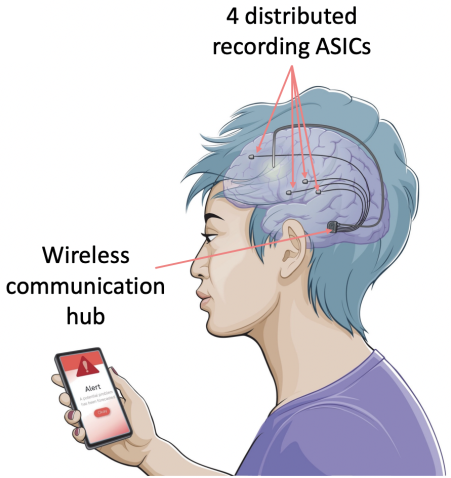

Real-time seizure detection

Talk by W. Lemaire