Fermilab

Modeling and Adaptive ML Based Control for SC Magnet Systems

Maira Khan, Abhijith Gadrankota, Stoyan Stoynev, Vittorio Marinozzi, Nhan Tran, Jennifer Ngadiuba, Rose Powers, Isobel Ojalvo

Contact Info: maira@fnal.gov

29 June 2026

U.S. DEPARTMENT of ENERGY

Fermi National Accelerator Laboratory is managed by FermiForward for the U.S. Department of Energy Office of Science

FERMI NATIONAL ACCELERATOR LABORATORY

Outline

1. Introduction
2. Magnet and RF Needs at the Muon Collider
3. HTS Magnets Introduction
4. HTS Challenges
5. Previous AI R&D Efforts at FNAL on Surrogate Modeling for HTS Design
6. Future Outlook

Introduction: Livingston Plot

From V. Shiltsev

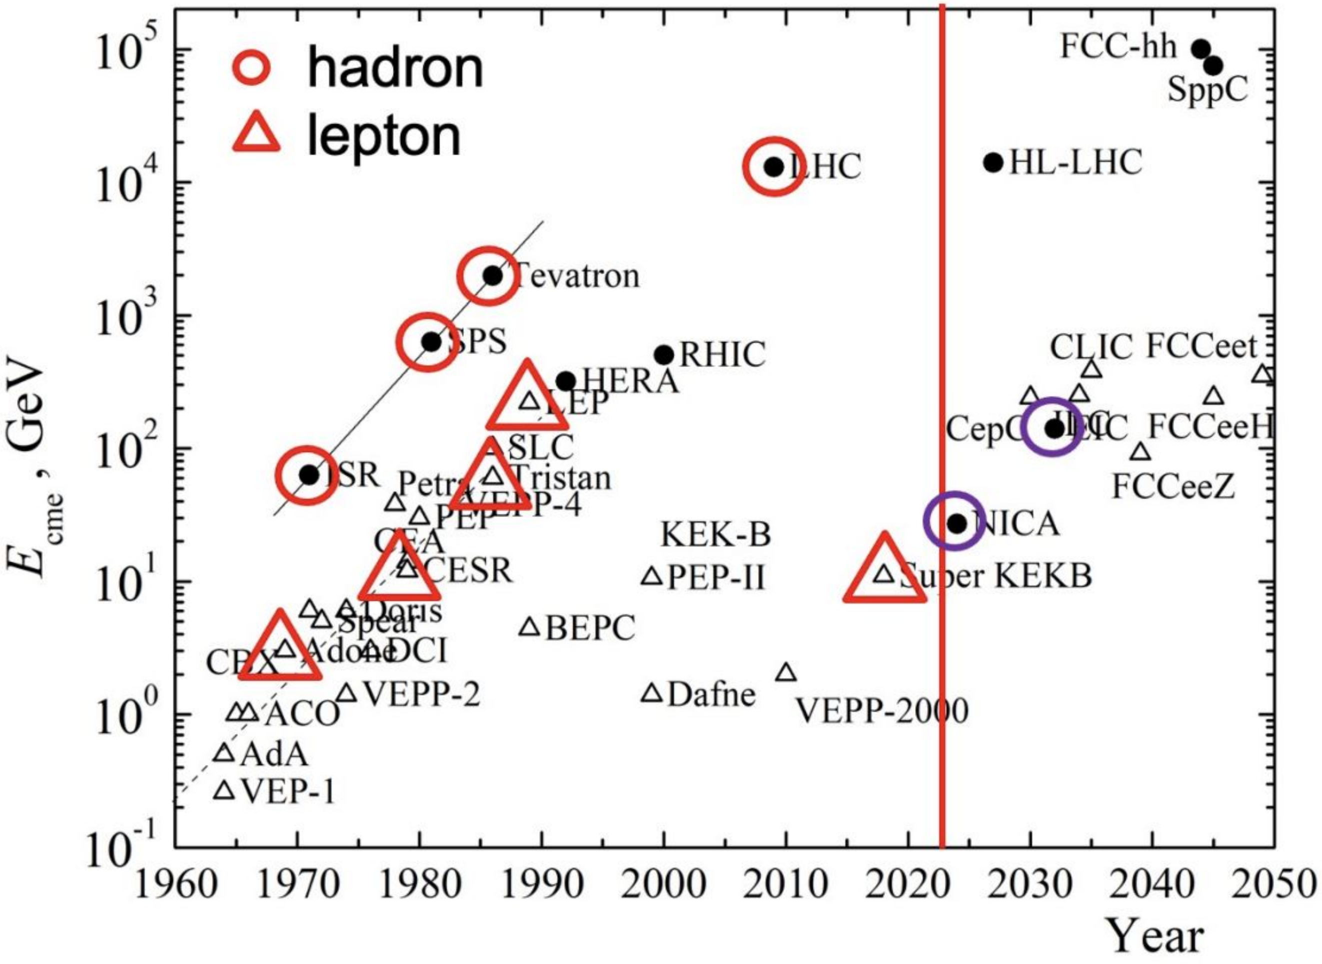

- **The Livingston plot illustrates that energy-frontier progress is driven by successive technology breakthroughs.** Each generation of colliders eventually reaches practical limits, requiring new accelerator technologies to continue the exponential growth in collision energy.
- **High-field superconducting dipoles are the next enabling technology.** Low synchrotron radiation allows multi-TeV muon beams to be bent with **Nb_3Sn (~11 T)** and future **HTS (14–20 T)** magnets, making a compact energy-frontier collider possible.
- **Ultra-high-field HTS solenoids enable ionization cooling.** Rapid cooling before muon decay requires **20-30+ T** solenoids to produce the low-emittance, high-brightness beams needed for high luminosity.
- **Rapid-cycling magnets enable acceleration within the muon lifetime.** Fast-ramping normal-conducting magnets, together with fixed superconducting magnets, accelerate muons to collision energy before they decay.

Magnet and RF Needs

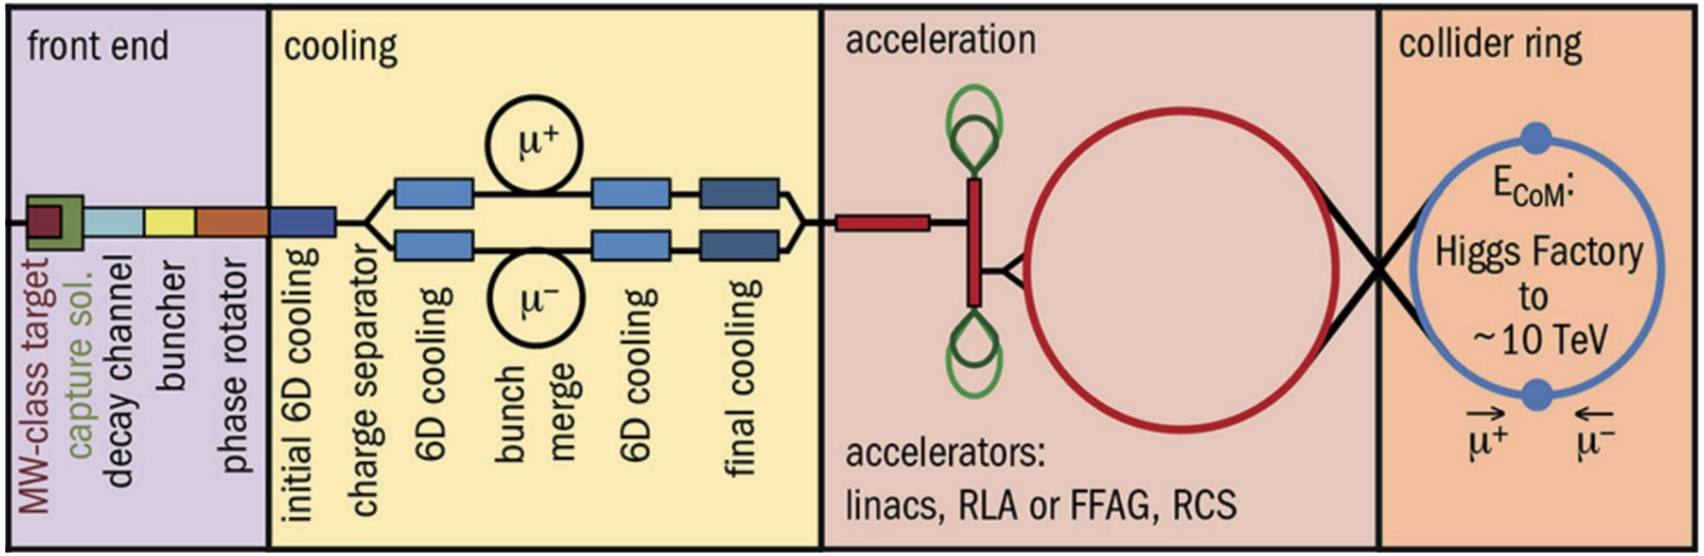

Schematic layout of the proton-driven muon collider. Source: L. Bottura, CERN Yellow Reports CERN-2024-003 (reproduced from R. Palmer); IMCC.

Stage	Superconducting Magnets	RF Role
Front End	High-field capture solenoids collect pions and muons; transport solenoids guide the beam.	Bunches the beam and reduces energy spread (buncher & phase rotator).
Cooling	Solenoids High Field (>20T?!) to focus the beam for ionization cooling	Restores only the longitudinal momentum lost in absorbers, enabling emittance reduction.
Acceleration	Dipoles and quadrupoles guide and focus the beam; hybrid rapid-cycling magnets enable fast acceleration.	Provides the accelerating gradient needed to reach TeV energies before muon decay.
Collider Ring	High-field dipoles determine collider size; final-focus quadrupoles maximize luminosity.	RF maintains beam energy and compensates for small energy losses during operation.

HTS Magnets for the Muon Collider

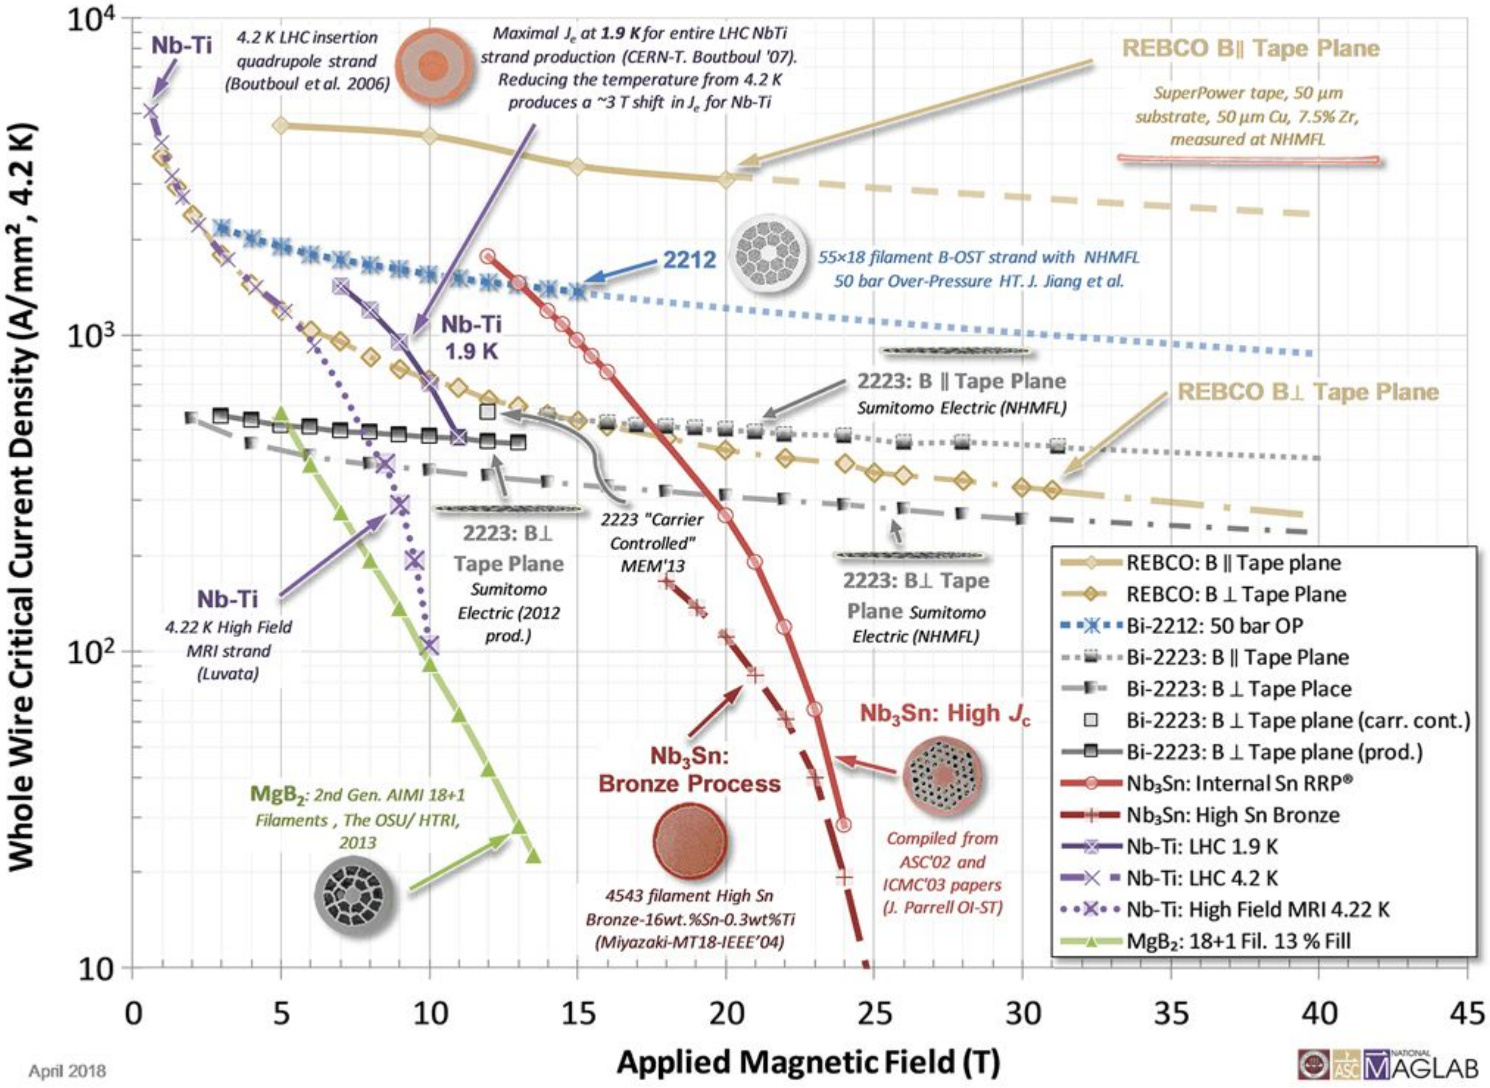

Figure: National MagLab at FSU (2018)

- HTS conductor shows high promise for higher field and current density to achieve the necessary field strength for muon collider needs
- Non-adiabatic quench allows us to re-ramp current during O(10ms)

HTS Magnet Challenges

- *If HTS is so promising, why aren't we there yet?*
  - HTS arrives as a flat, mostly-not-superconductor tape — awkward to cable, bends about only one axis, and drives the magnetization/AC-loss problems.
  - Critical current depends strongly on field *angle* (factor 3–5); field orientation varies around a coil, so you design for the worst local angle
  - Above ~40–50 T the magnetic pressure rivals material strength — the final-cooling solenoid becomes a structures problem (≈650 MPa hoop, no tension allowed, even in quench).
  - Quench Protection: HTS is *too* stable — the normal zone barely propagates, so a quench is hard to detect before energy concentrates and burns a hot spot.

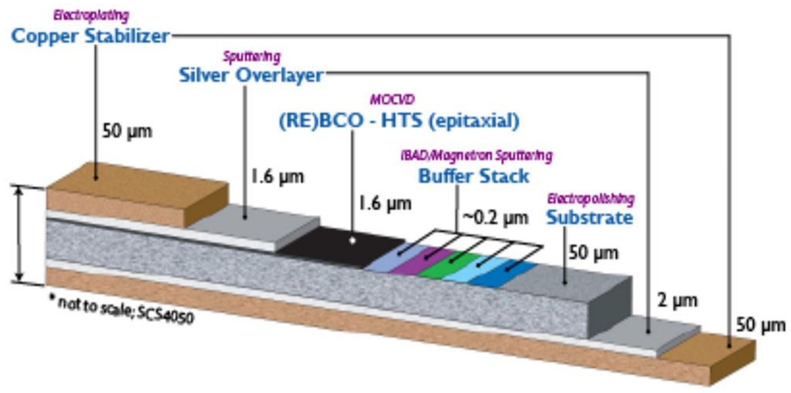

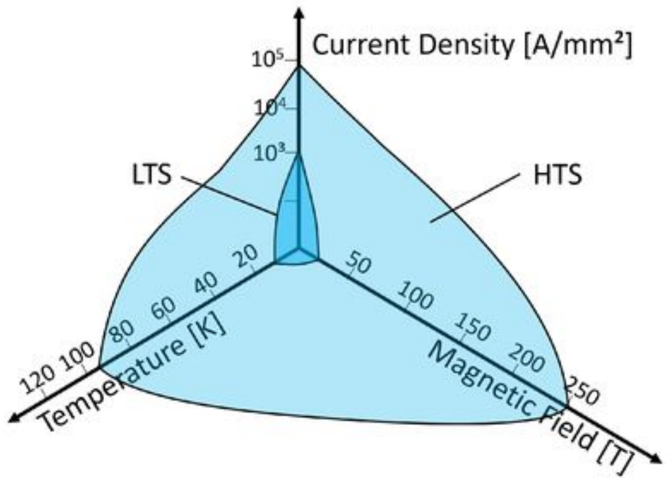

AI in the Loop?

A few demonstrated cases….

Case 1: Early Quench Protection from Data Alone

Under AIEHEP Project (Ngadiuba, Tran) and LDRD (Gandrakota et. al) with FNAL Magnets Department we have been able to use autoencoders, weakly supervised learning, and anomaly detection approaches to stream signals and identify events associated to quench with some model based probability. We are still waiting on HTS data to deploy diagnostics

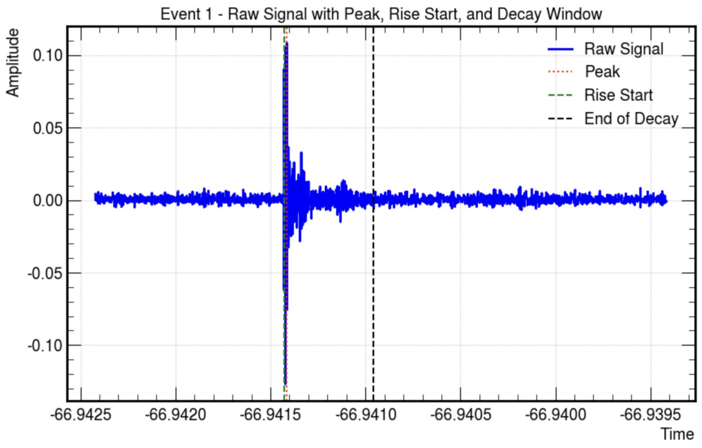

Instrumentation Processing

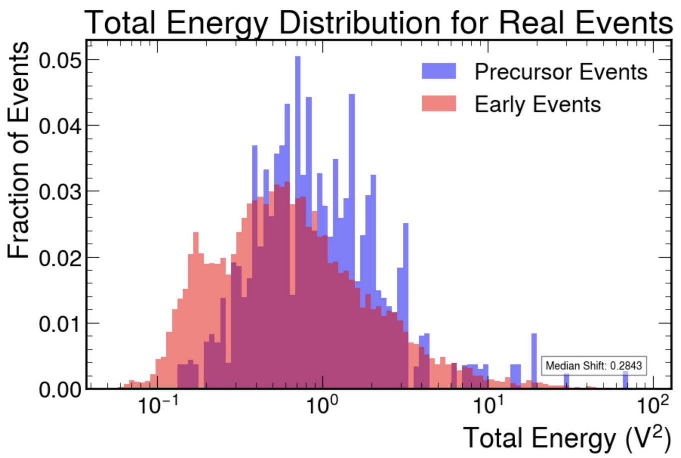

Analysis

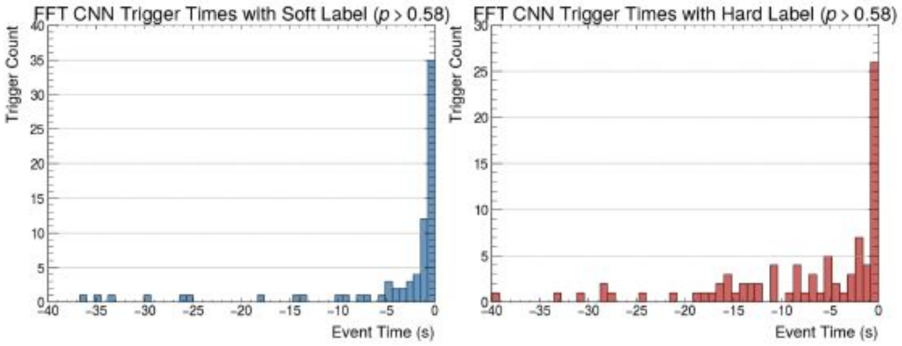

Real Time Trigger

Fig. 1-3: Event selection, analysis and trigger procedure as described in [paper]

FERMI NATIONAL ACCELERATOR LABORATORY

Case 2: Surrogate and RL Control in the Loop

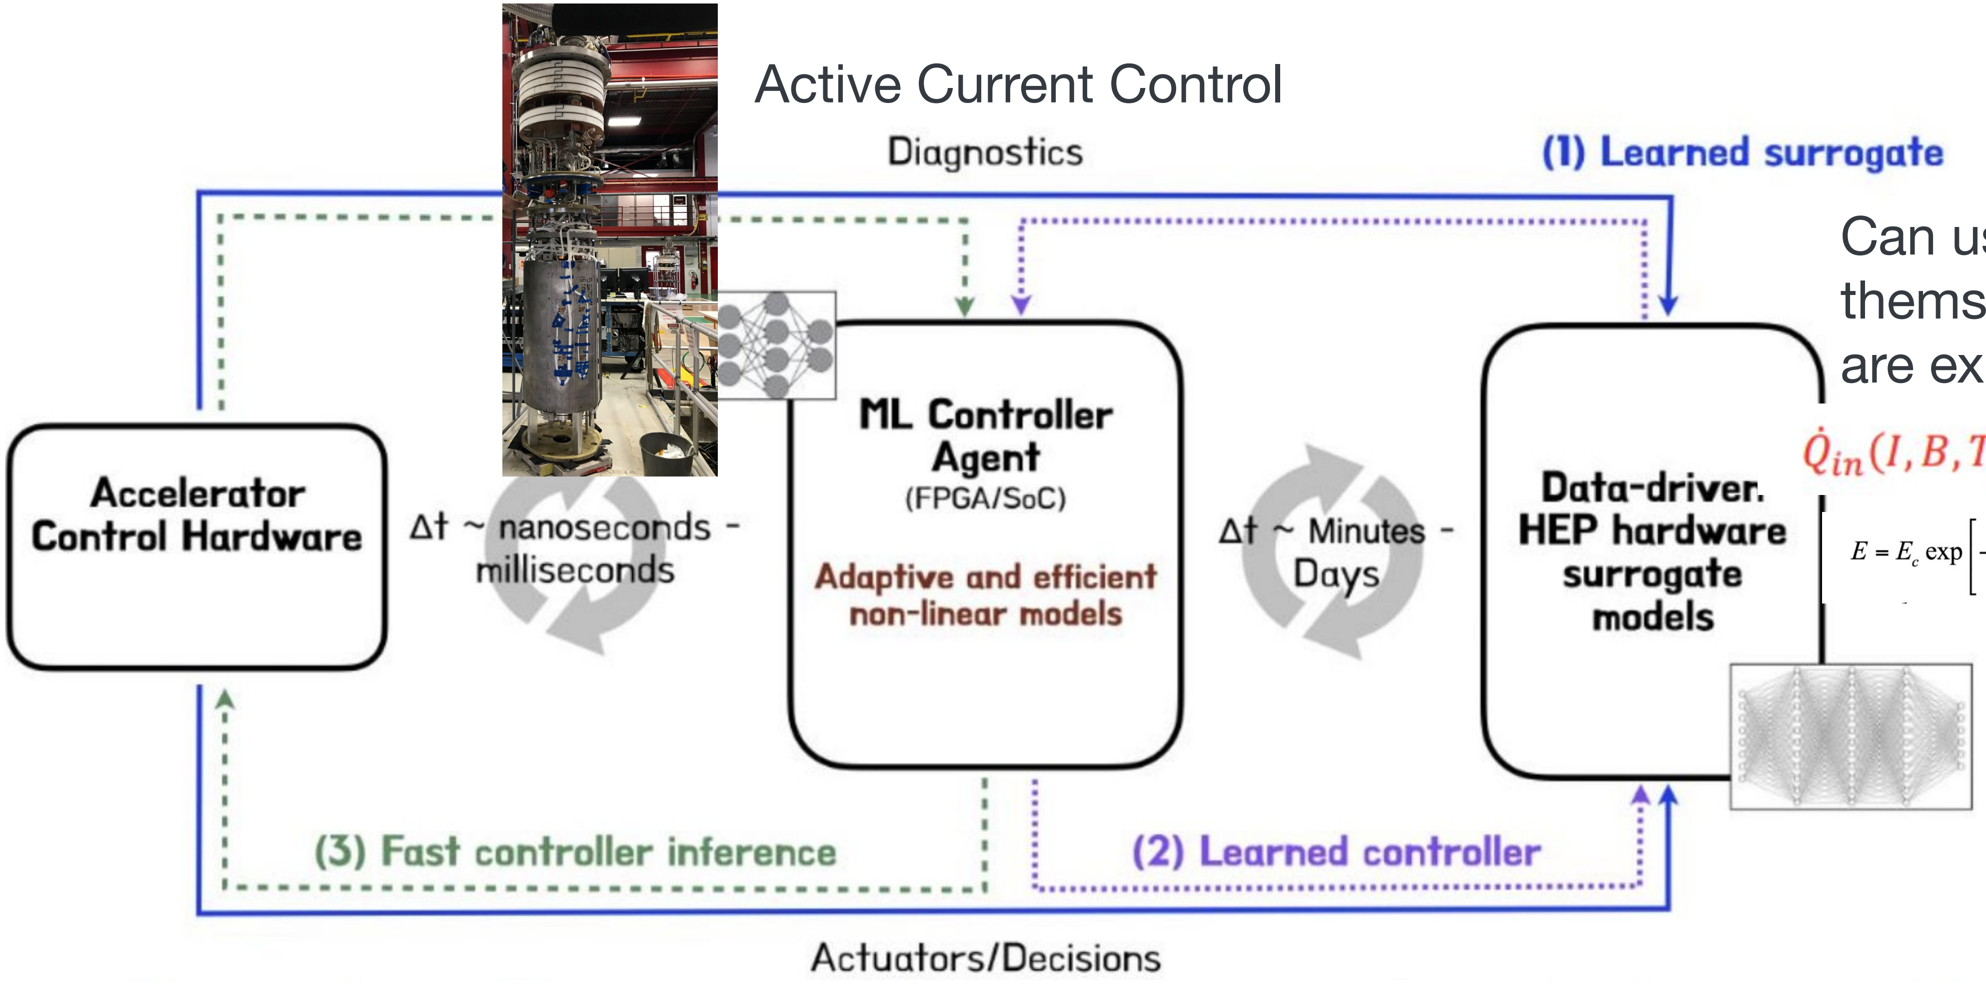

Can use gradients themselves as they are expressive!

$$\dot{Q}_{in}(I,B,T) = \dot{Q}_{out}(T)$$

$$E = E_c \exp\left[-n(H,T)\left(1-\frac{J}{J_c(H,T)}\right)\right]$$

Thermal Runaway Approach

- M. Marchevsky and S. Prestemon proposed and experimentally validated a thermal runaway criterion for HTS conductors published under *M Marchevsky and S Prestemon 2024 Supercond. Sci. Technol. 37 035012*
- HTS conductors may sustain a stable dissipative flux-flow regime prior to irreversible thermal runaway
- Thermal stability is governed by the balance between Joule heating and cooling
- Thermal runaway occurs when no stable equilibrium solution exists

A condition for the dynamic equilibrium: $\dot{Q}_{in}(I,B,T) = \dot{Q}_{out}(T)$

$$\dot{Q}_{in} = I^2 R(I,B,T)$$

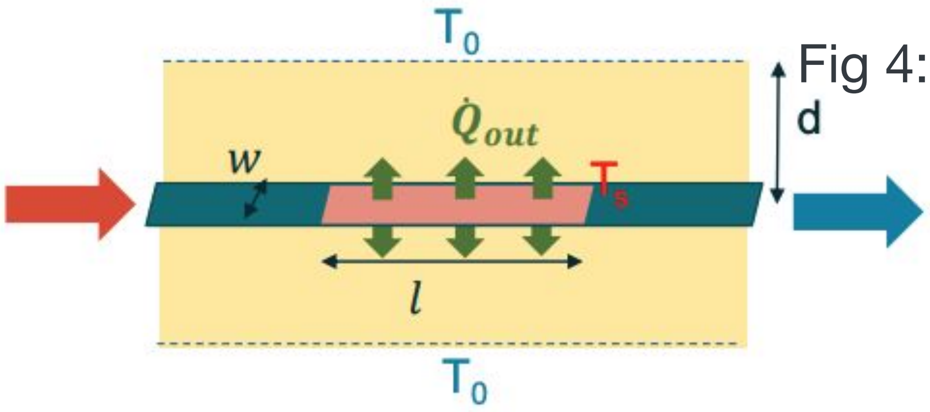

Fig 4: Credit to M. Marchevsky

$$\dot{Q}_{out} = \lambda(T)\frac{2lw}{d}(T - T_0)$$

$$R_s = E_0 l \frac{I^{n-1}}{I_c(B,T)^n}$$

As long as this condition can be satisfied for some current I^* (and the corresponding $T^* > T_0$), the flux flow regime is **stable, and there is no thermal runaway**.

Reproducing the Analytic Solver

- Using the thermal runaway framework described in the paper we can numerically generate operational stability manifolds across a range of conductor and cooling conditions for stabilized conductor surface. See backup for cross check of surfaces
- Higher cooling strength increases the stable operating current before runaway
- Increasing base temperature reduces the available thermal stability margin

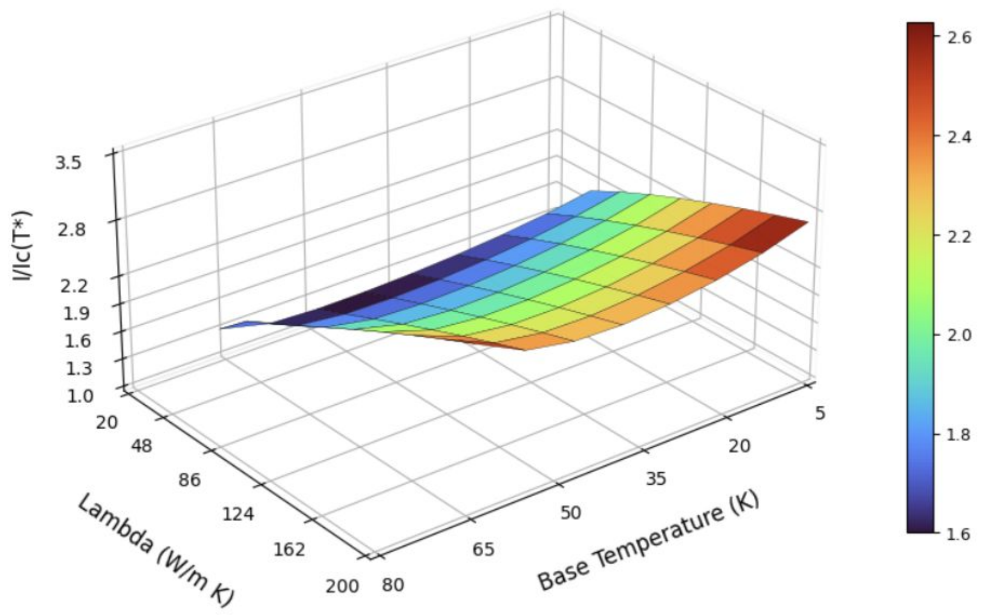

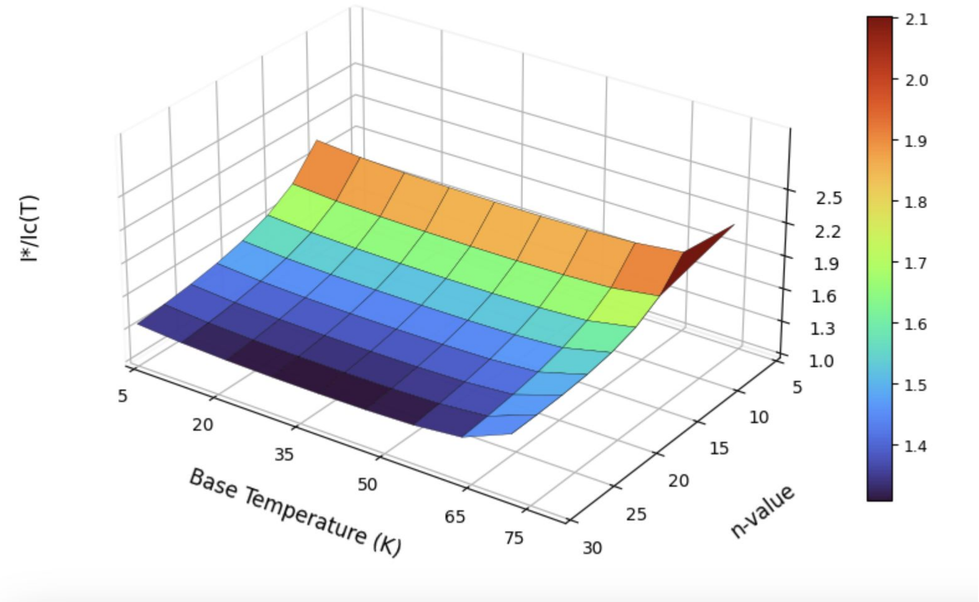

- Left: runaway current ratio versus base temperature and cooling coefficient λ
- Right: Runaway current ratio I*/Ic(T) versus base temperature T_0 and n-value of the power-law E–J relation."

Surface Reproduction: Analytic vs Surrogates

- Physical behavior captured by both surrogates
  - Broad plateau at low T_0 where cooling dominates
  - Up-turn at high T_0 (>60 K) as current redistributes into the stabilizer
  - Saddle-shape geometry preserved by both surrogates

Hoping to get to FPGA implementation soon!

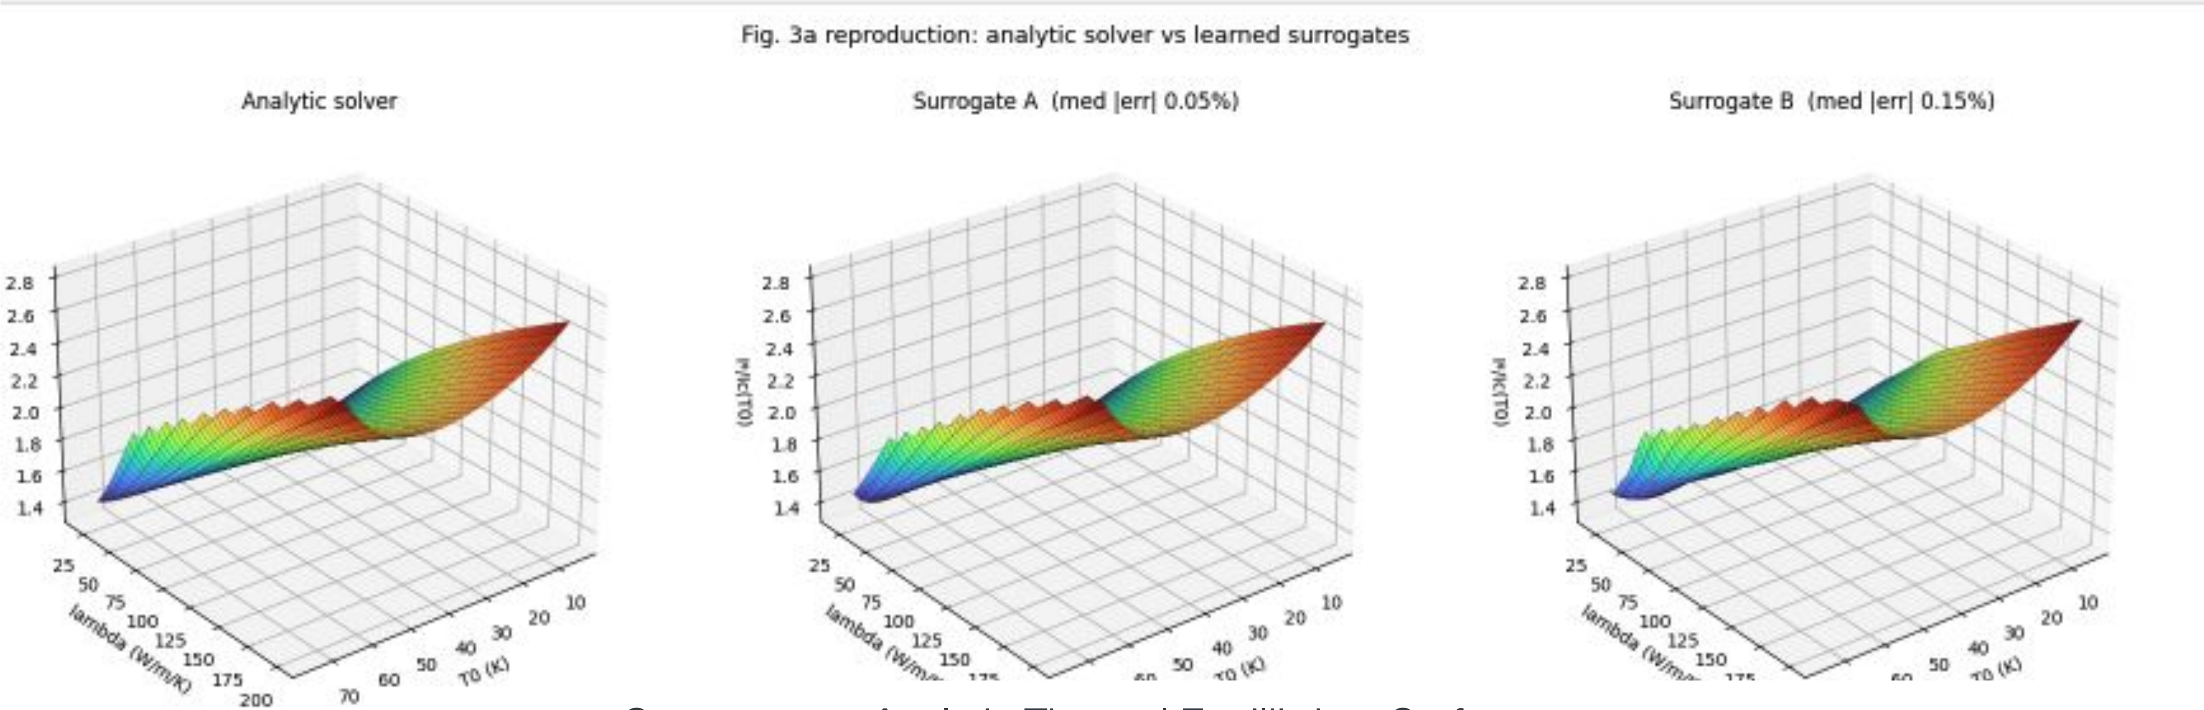

Surrogate vs. Analytic Thermal Equilibrium Surfaces

Data Taking Campaign This Summer for HTS

Future: Training Control Policies

In Progress: Integrated Thermal and Electromagnetic Surrogate with Calibrated Instrumentation Response

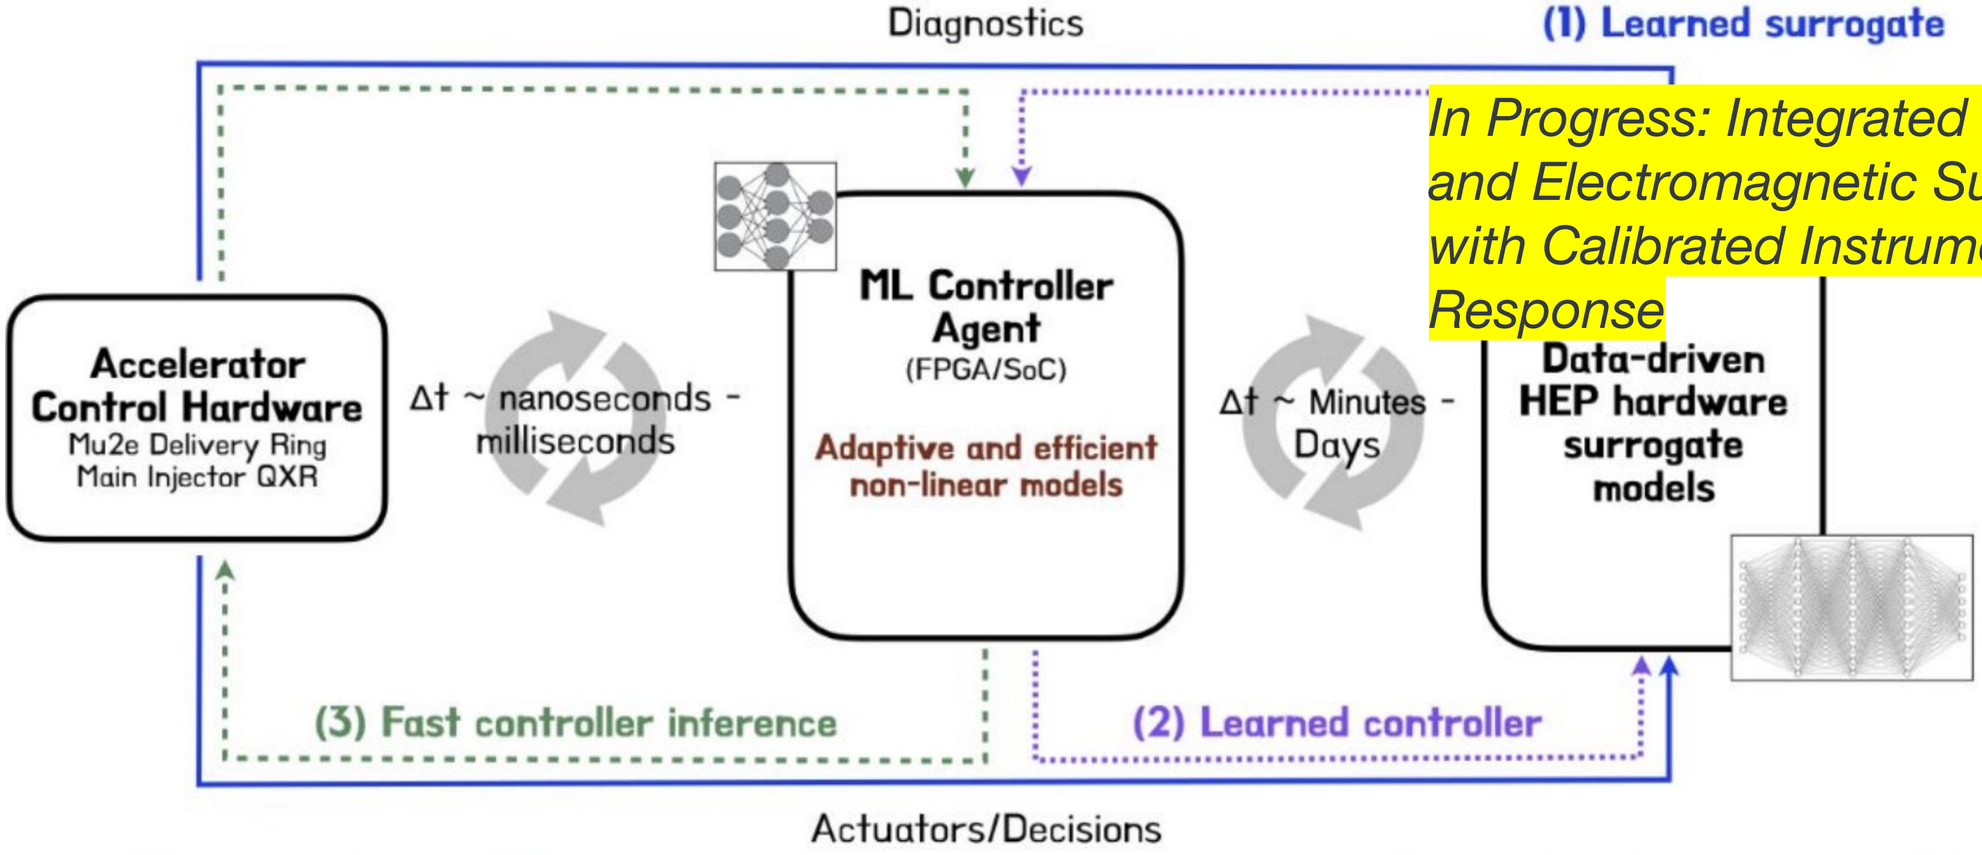

An Example: Mu2e Extraction RL Control with Surrogate

Differences between Two RL Controllers

Train in surrogate

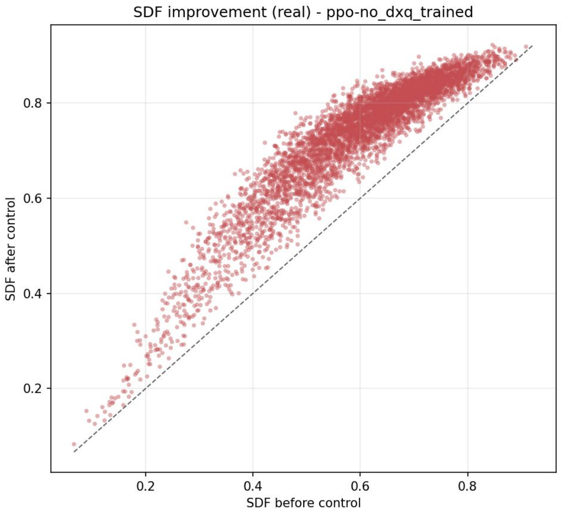

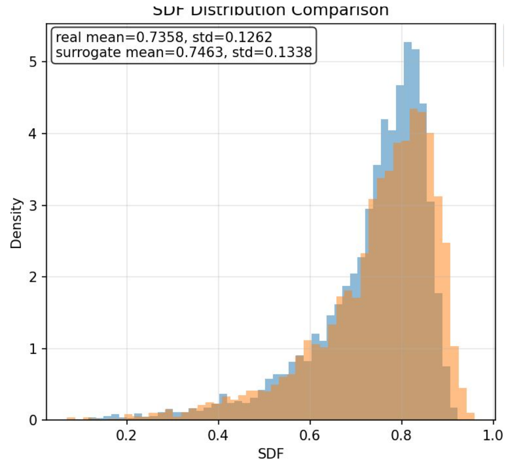

Train in sim

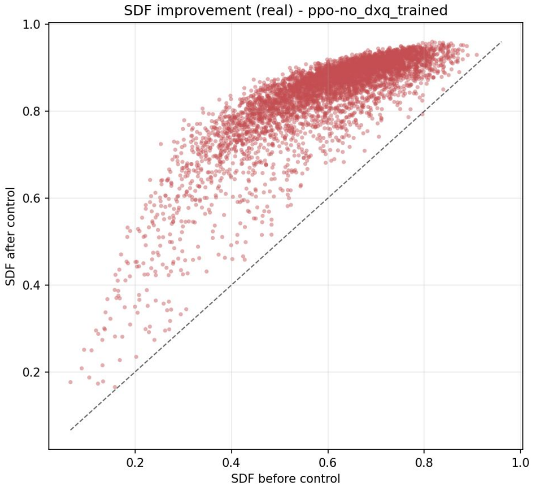

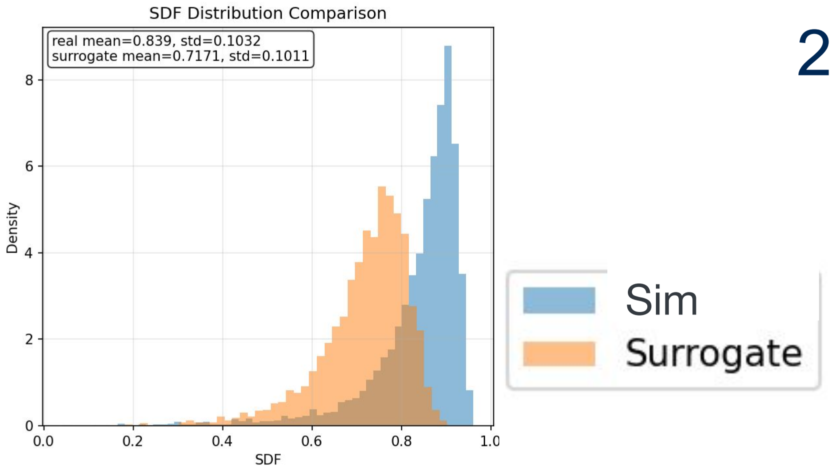

RL Controller (0.839)

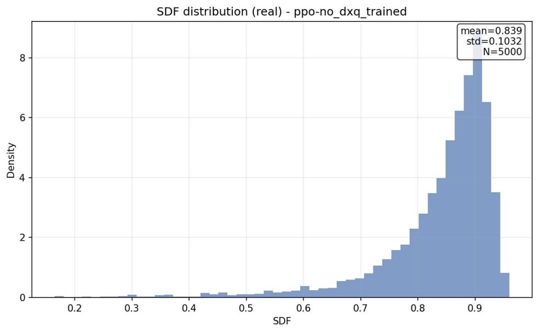

PID Controller (0.645)

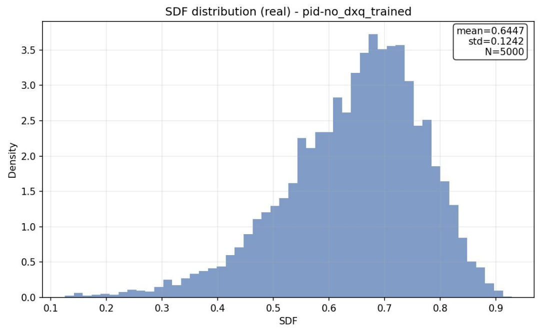

1. RL controller trained on simulator has more large actions
2. Hypothesis: Large Actions may be necessary for the case of low-SDF spills

Similar efforts going on for resonance control at FNAL also under A. Whitbeck et. al

Thoughts for Future

AI for Digital Twins: Single Tape to Coil Scalability

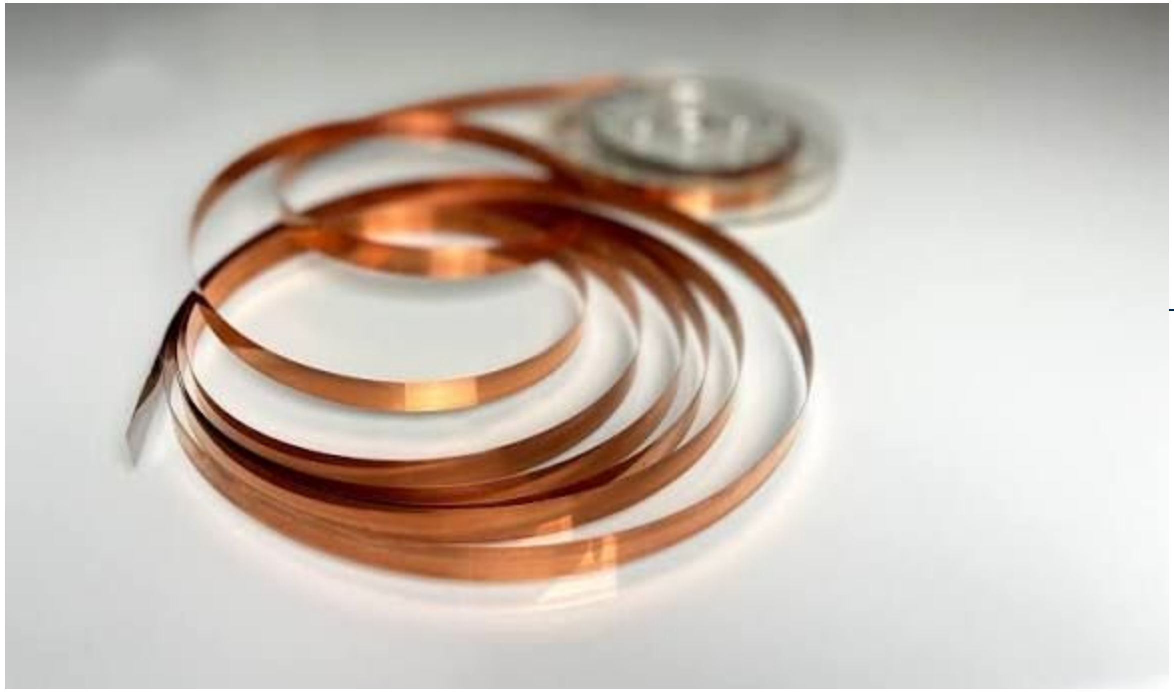

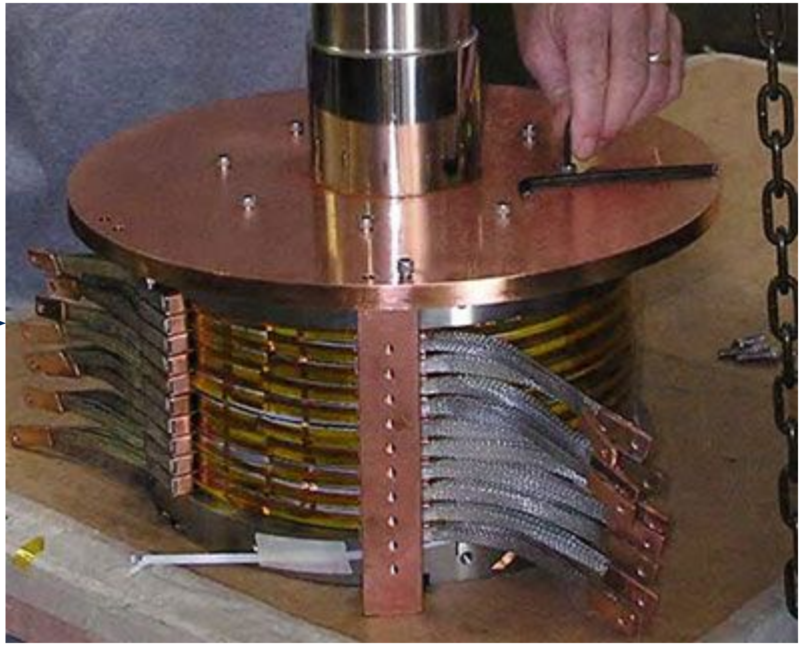

How to model optimal coil winding and configurations with AI? A digital twin sandbox? Eventual agentic DAQ integration?

Options: Fast FEM Surrogates, PINN Based, Inverse Reconstruction from Circuit Constraints

AI for Top Down Design

Turn the $J_c(B, \theta)$ field-dependence problem into an inverse-design loop?

The design loop:

- **Target spec** — start from what you want: bore field, aperture, field quality, and stress / margin / cost limits
- **Generate & optimize** — propose geometries via topology optimization or generative models; search with multi-objective Bayesian optimization
- **Surrogate multiphysics** — neural-net surrogates predict field, stress, and quench margin orders of magnitude faster than FEM, with $J_c(B, \theta, T)$ embedded as a hard constraint
- **Validate in FEM** — run full multiphysics only on the Pareto-optimal candidates, then iterate

In Parallel: AI for Resonance Control *(Whitbeck. et al)* at FNAL

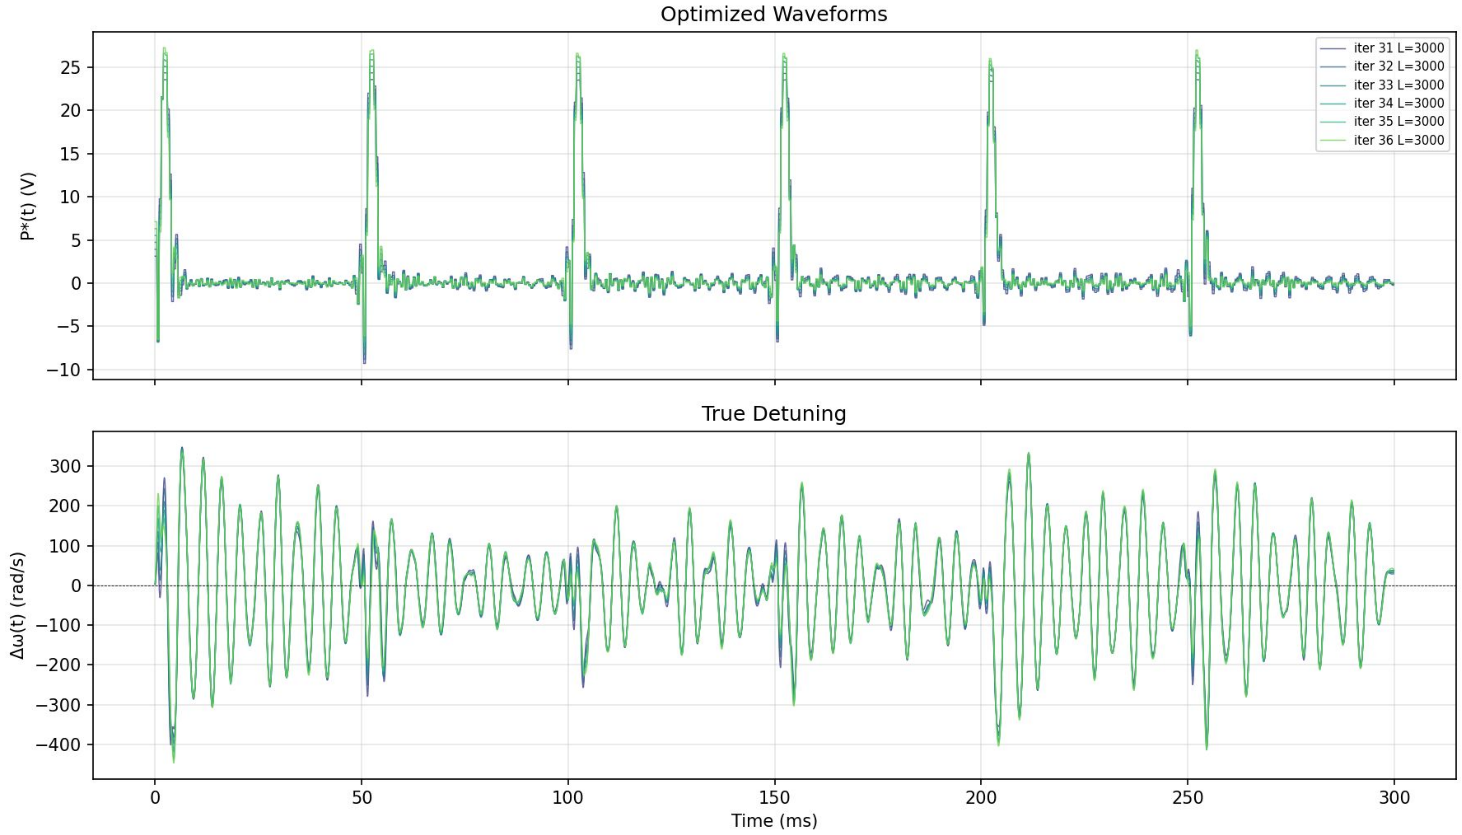

- Well defined analytic problem
- Disturbance shown (≈ ±64 Hz) exceeds spec for PIP-II ~2–3×: the detuning-dominated regime active control must tame
- Converged feedforward (iter 31–36 coincident) cancels the deterministic, pulse-synchronous LFD (Lorentz Force Detuning) transient — the dominant 20 Hz contribution
- Less power in pulsed mode, though more challenging!

Let's discuss!