Image credit: Marguerite Tonjes

Differentiable Simulation for Muon Cooling

AI4MuC mini-workshop

Sergo Jindariani, Nikita Kuklev, Prasanth Shyamsundar, **Nick Smith**, Diktys Stratakis, Nhan Tran (FNAL)
Philip Harris, Andrzej Novak, Holger Witte (MIT)
Eric Anderson, Corrinne Mills (UIC)

June 29, 2026

The challenge

- A muon collider at FNAL would be cool, but the beam must be cooled
- Rectilinear design: 10^5 reduction in emittance through several stages in a 1km channel
  - Conservative design, hand-optimization
- Each of the 12 stages in the rectilinear ionization cooling channel has 20-100 cells, each of which has several tunable parameters
  - Optimal working point may differ in each stage, and improvements *compound*
- Simulation requires atypical mix of robust beam transport optics, material interaction, and collective beam dynamics

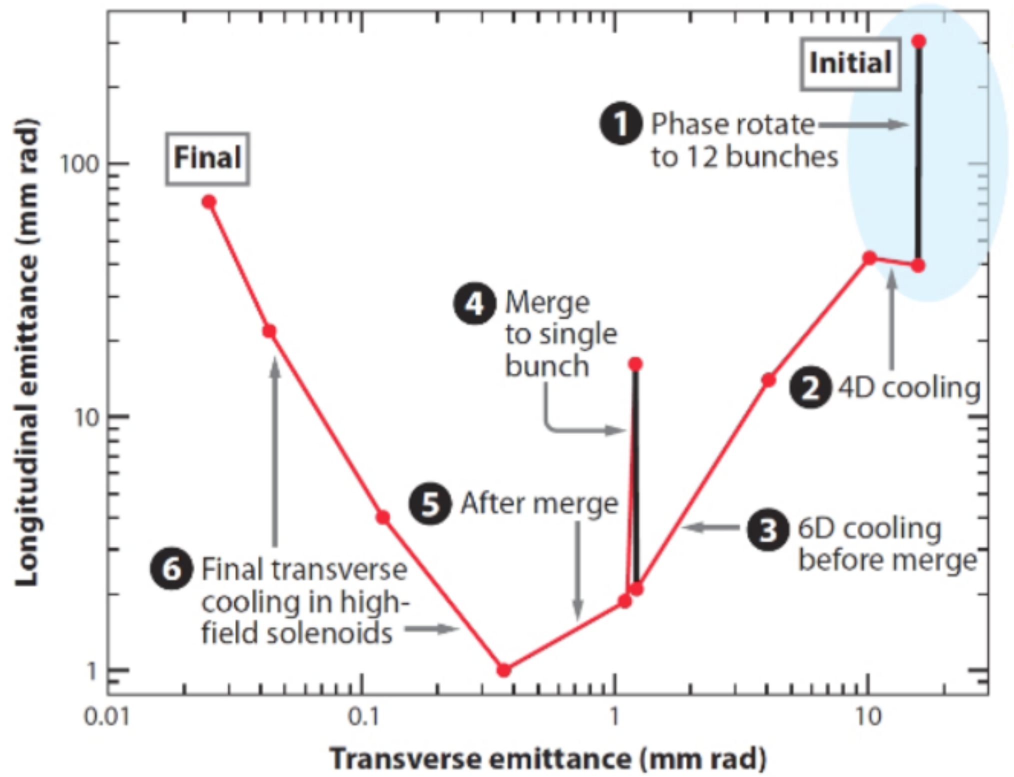

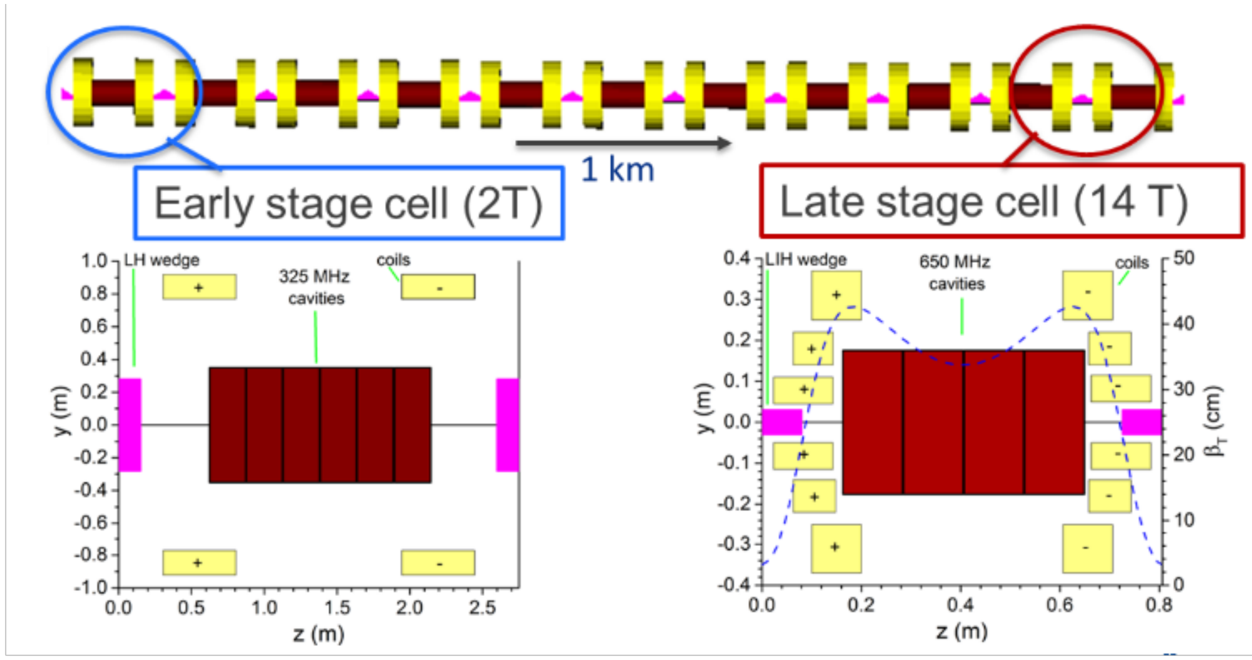

The opportunity

- Large community interested in applying differentiable programming techniques to HEP problems, such as detector design and beam tuning
  - See e.g. [MODE workshops]
- Differentiable simulation enables efficient high-dimensional optimization
  - For our problem, a reasonable set of building blocks is understood!
    - TM010 pillbox, thick solenoid, wedge absorber, … already a [benchmarking activity]

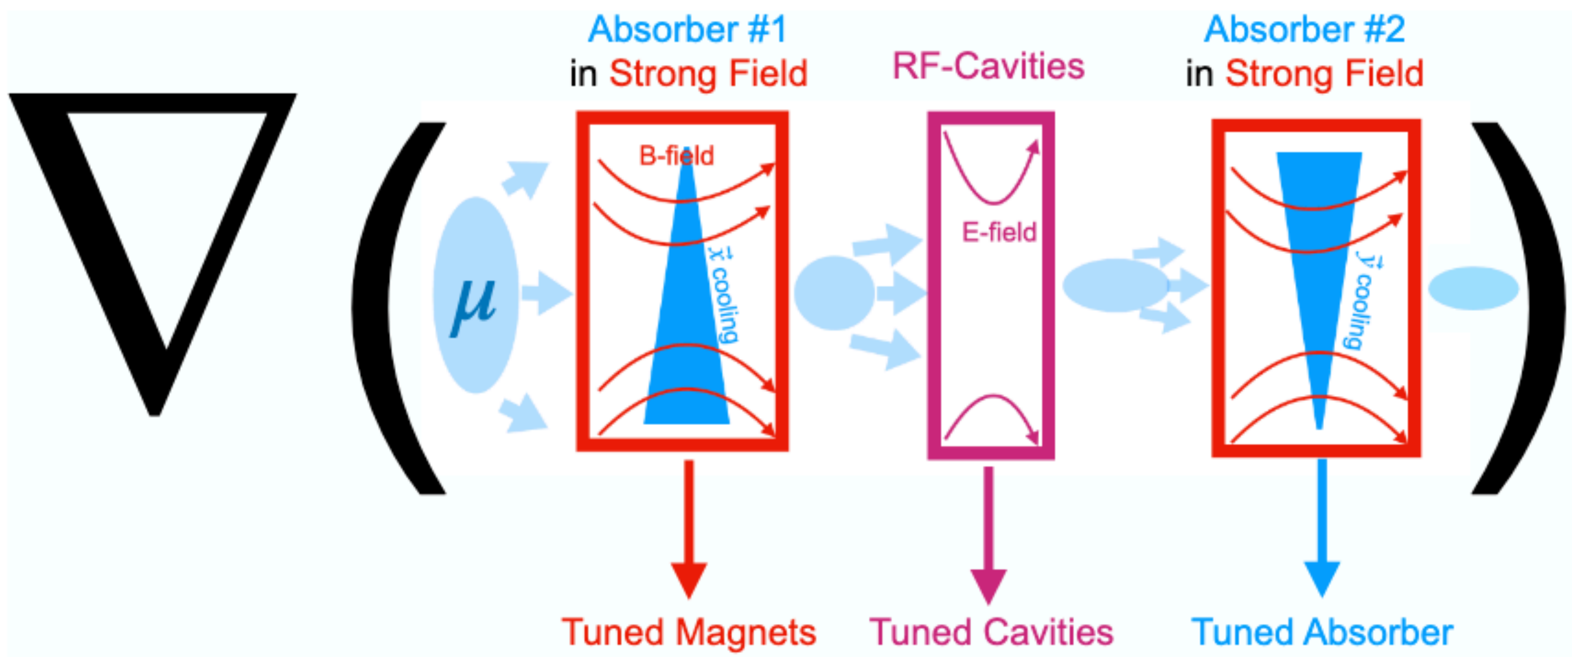

$$\frac{\partial \epsilon}{\partial B} = \frac{\partial \epsilon}{\partial A_2}\frac{\partial A_2}{\partial B} + \frac{\partial \epsilon}{\partial A_2}\frac{\partial A_2}{\partial RF}\frac{\partial RF}{\partial A_1}\frac{\partial A_1}{\partial B} + \dots$$

P. Harris

Fermilab

The plan

- There is plenty of prior art in both muon cooling simulators and differentiable beamline simulations to draw on
  - iCOOL, G4Beamline, BDSim, RF-Track, …
  - Cheetah, WarpX, Synergia, Bmad-Julia, …
- We are building a ground-up simulation in Jax
  - Lack of an existing tool at the intersection of all requirements
    - Optics + material interaction (+ collective effects) + ∂
  - Desire to leverage a computational framework with first-class (automatic*) differentiation support and GPU acceleration "for free”
  - **LLM coding tools change the calculus for refactor vs. rewrite**
    - At least for “Minimum Viable Product” stage
    - Once we understand the challenges, maybe we pivot?

* AD is not always correct nor always the fastest!

Progress

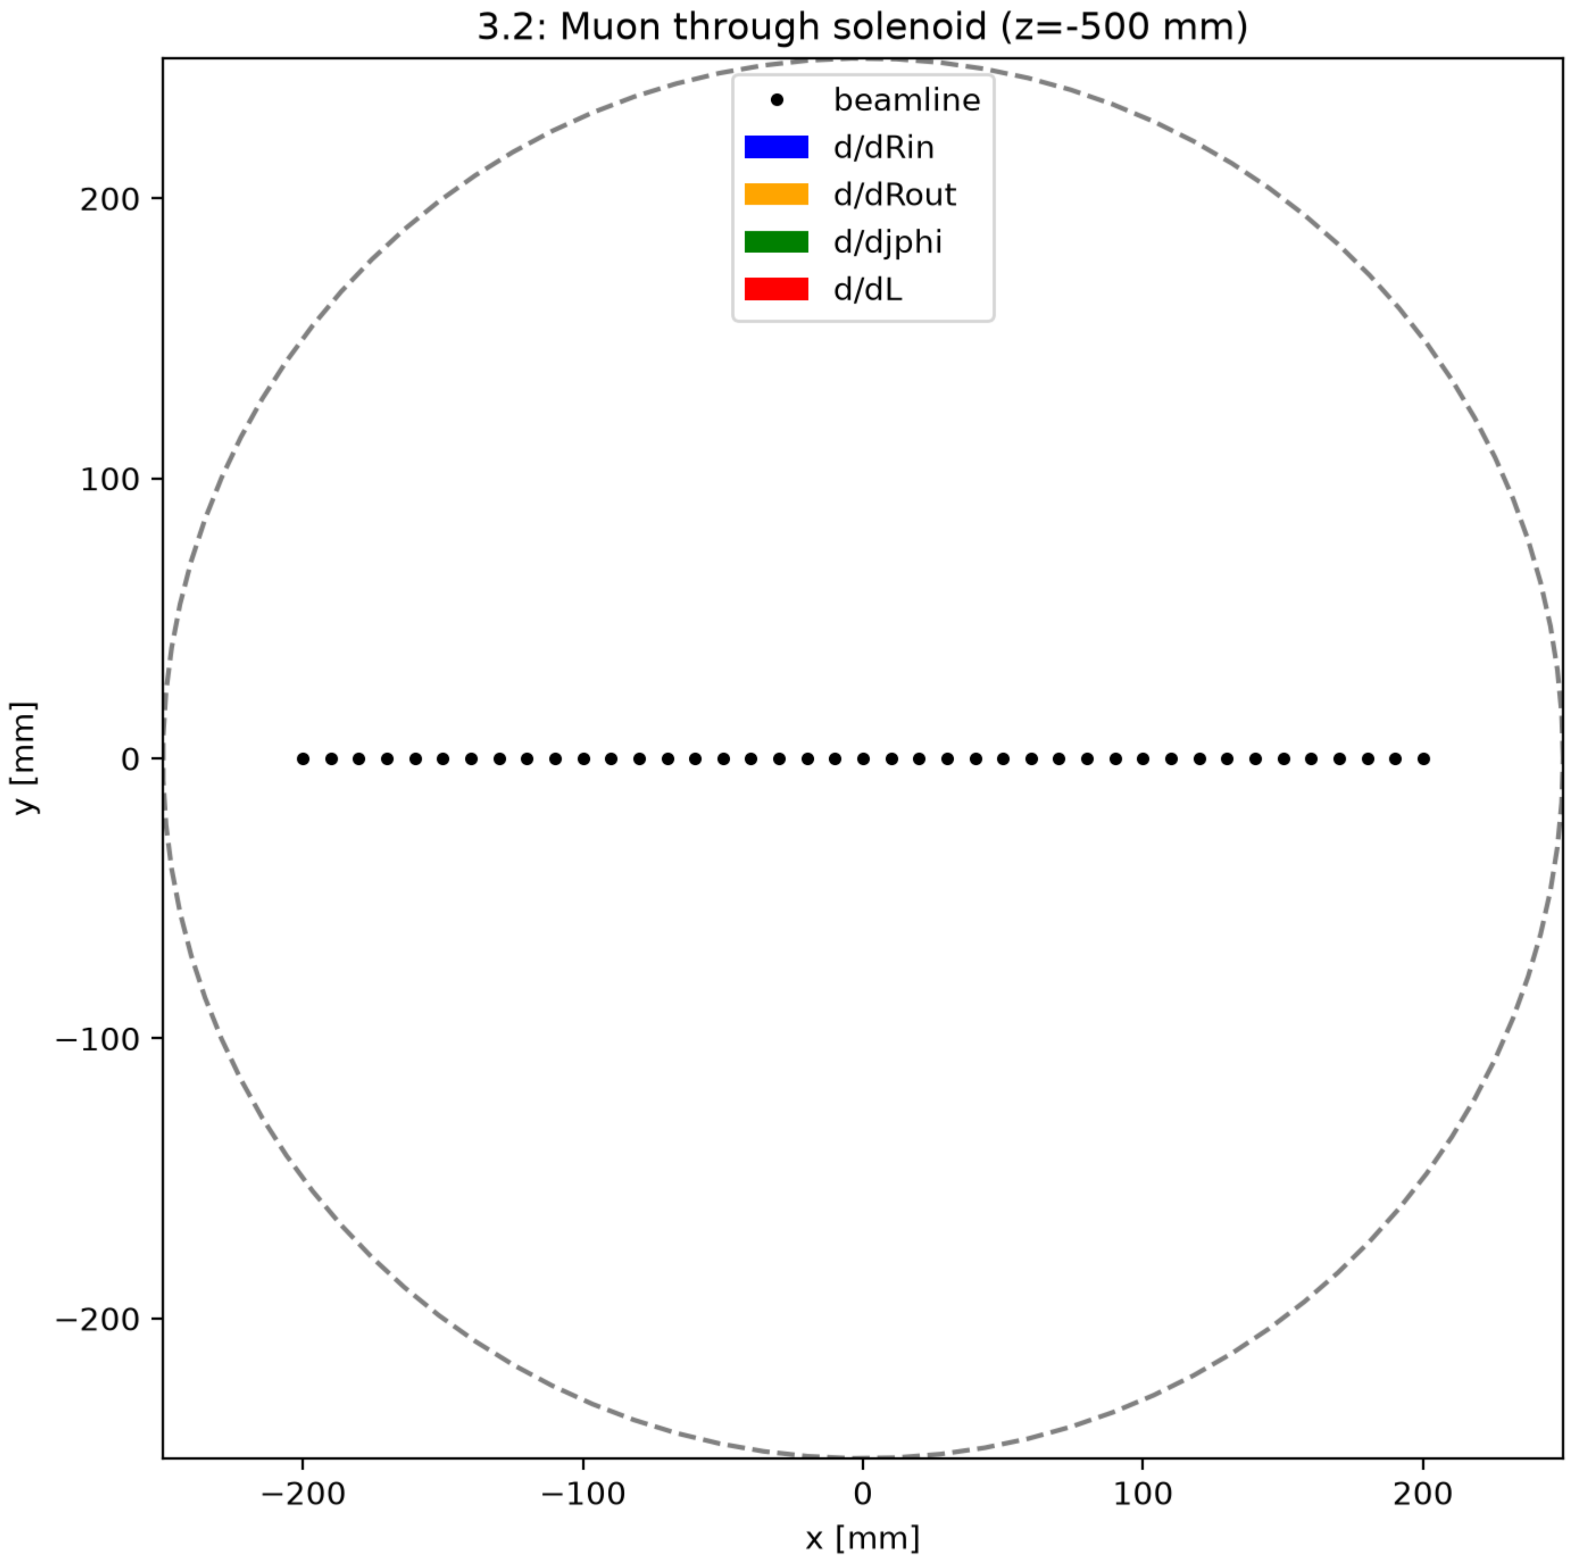

[Benchmark 3.2: thick solenoid]()

Progress

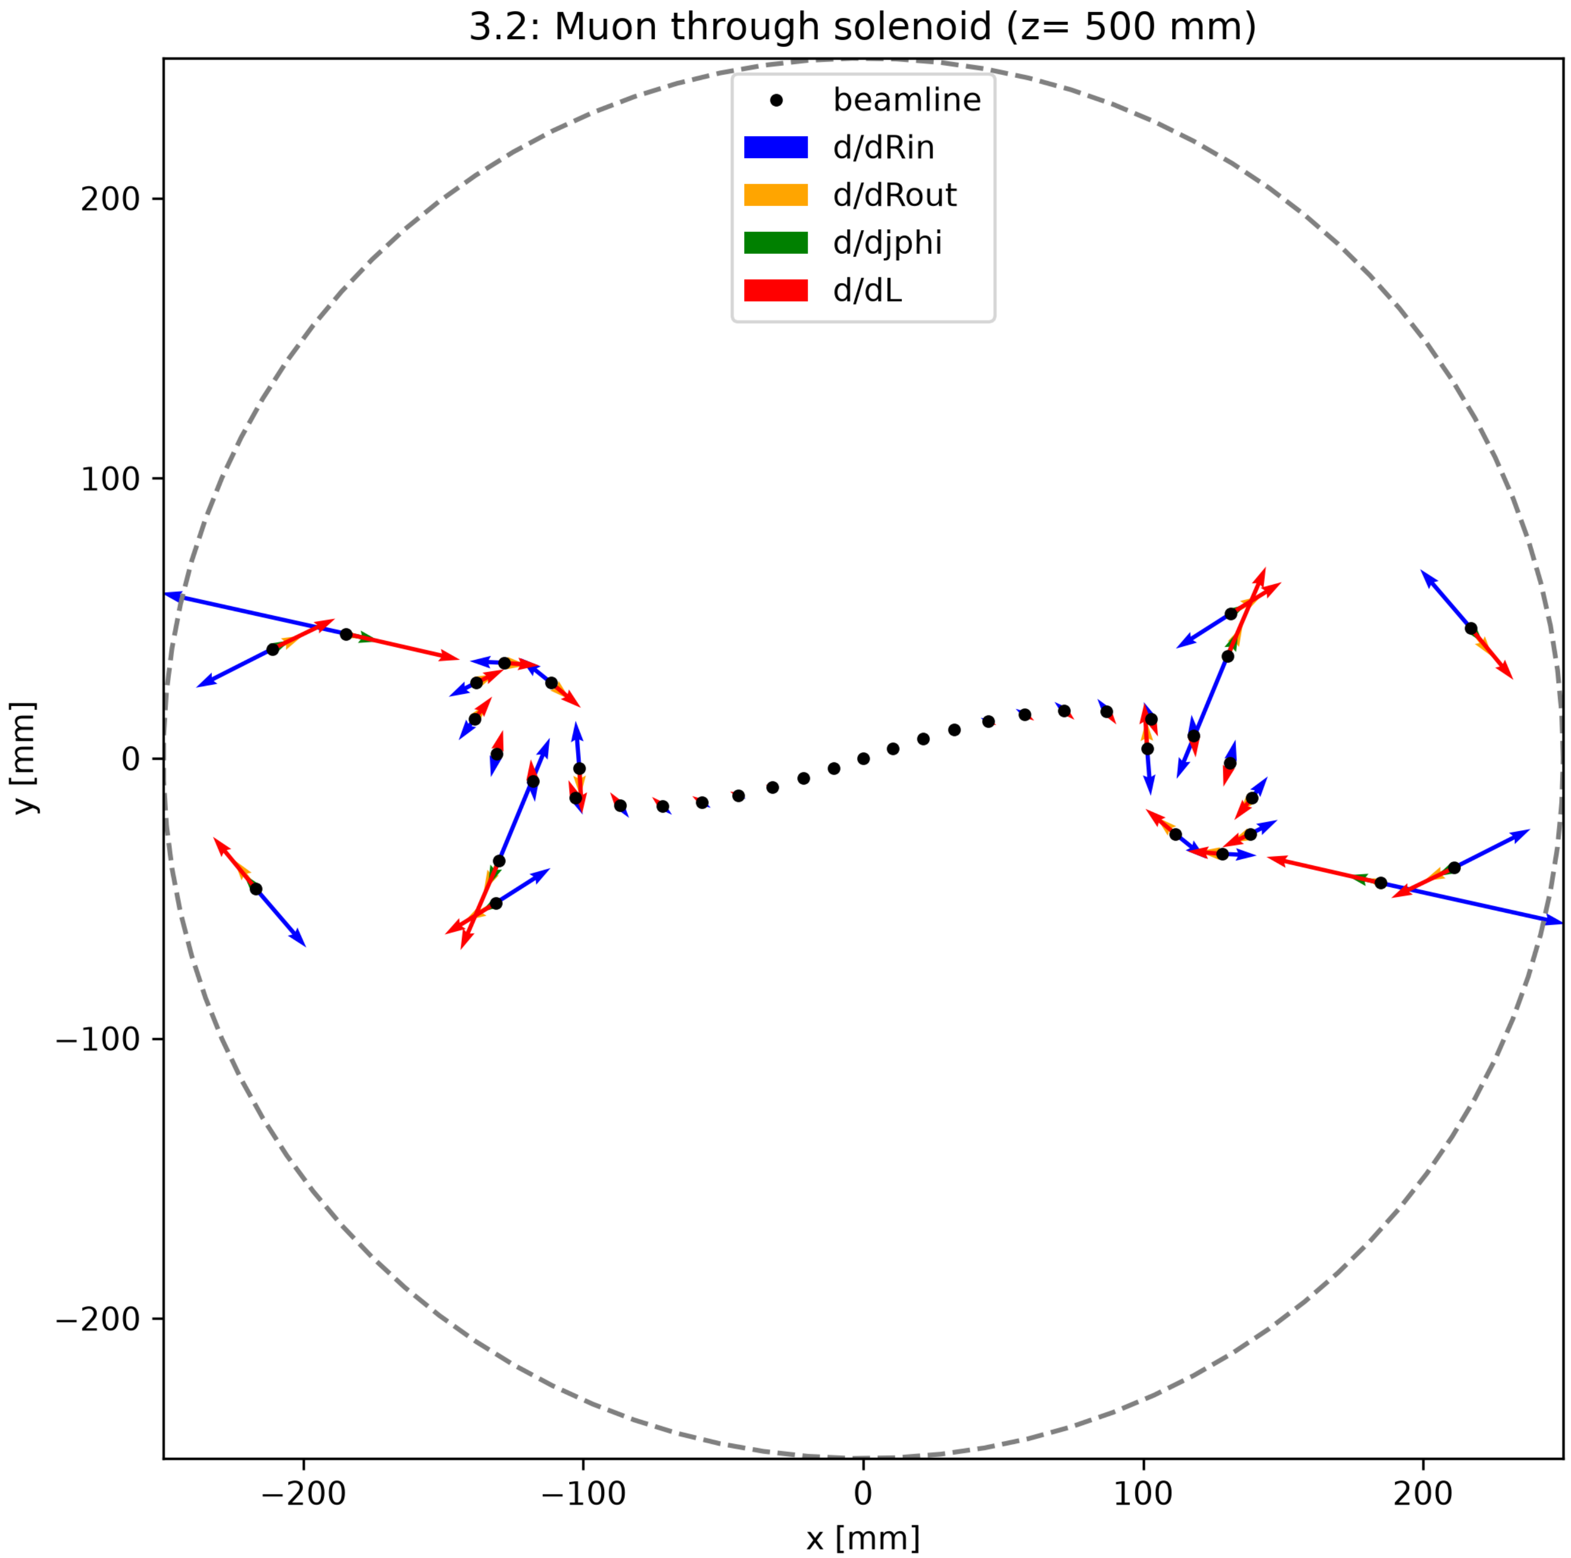

[Benchmark 3.2: thick solenoid]()

Stochastic effects

- The hard part is taking gradients through stochastic elements
  - Absorber: energy straggling, multiple scattering
- The objective is emittance, extracted from moments of 6D distribution $p(x, \dot{x})$
  - $\epsilon_\theta(t) \sim \int dx d\dot{x}\, f(x, \dot{x}) p_\theta(x, \dot{x}; t) = \int dx d\dot{x}\, f(x, \dot{x}) \int^t dt' \frac{dp_\theta(x, \dot{x})}{dt'}$
  - Sample $(x, \dot{x}) \sim p(x, \dot{x}; t = 0)$ and push forward $\Rightarrow$ weighted samples $(x_\theta, \dot{x}_\theta, w_\theta)$
  - Two paths for θ gradients: through the final positions, or through the sample weights
  - **The key for robust gradients is efficient sampling**

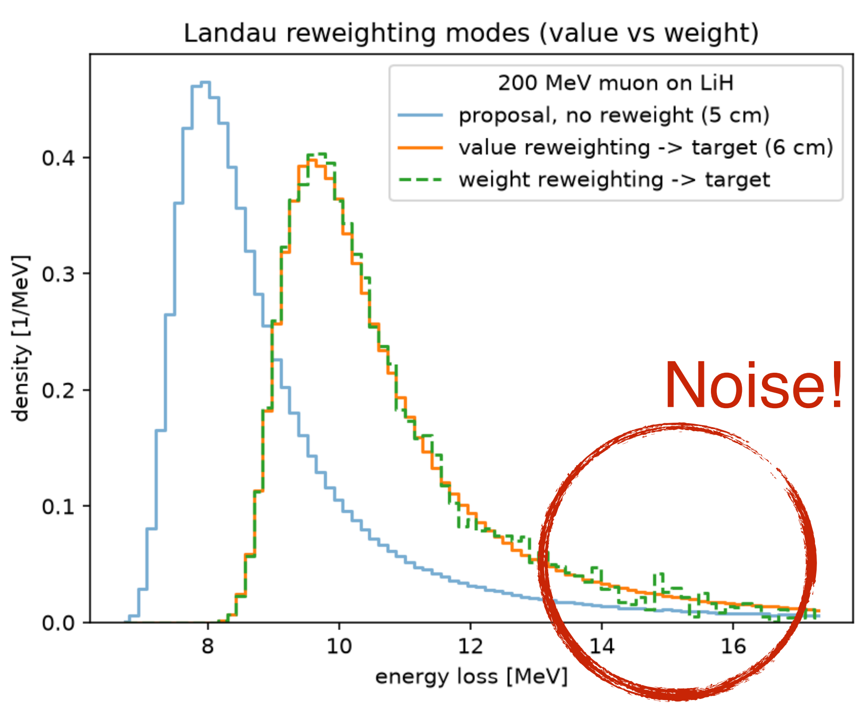

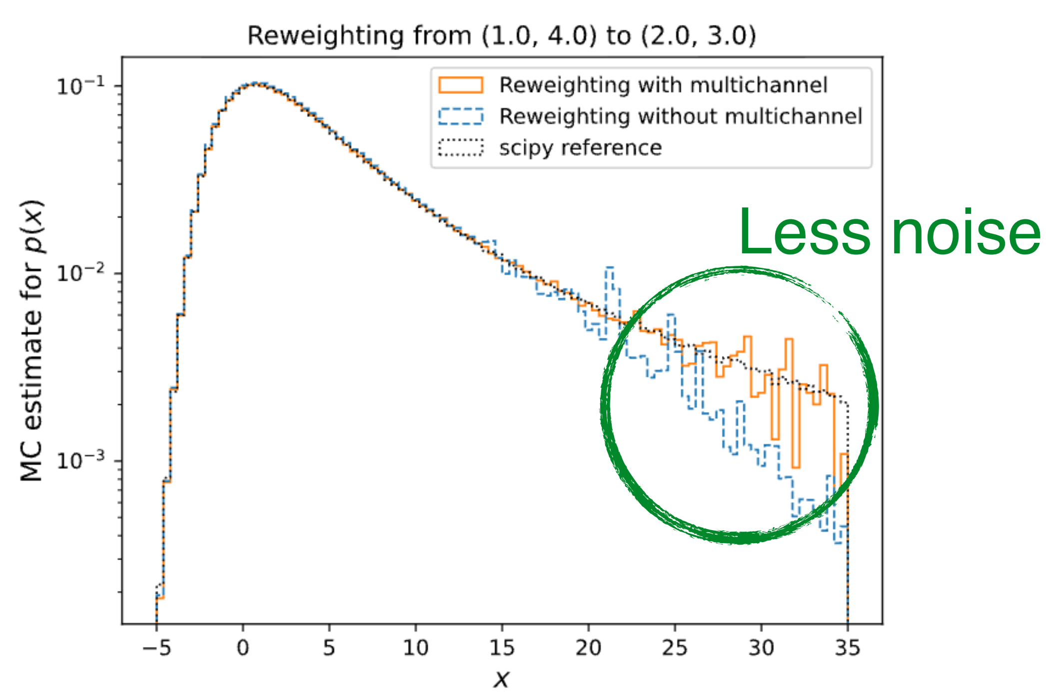

Future tie-ins

- A differentiable simulator becomes a very useful component of a digital twin for online beam controls
  - Easier to train NN to capture delta between simulation and reality than a black box!
  - To do so effectively, “classical component” needs high-dimensional gradients
  - Here we can tie in with [MOAT project]()

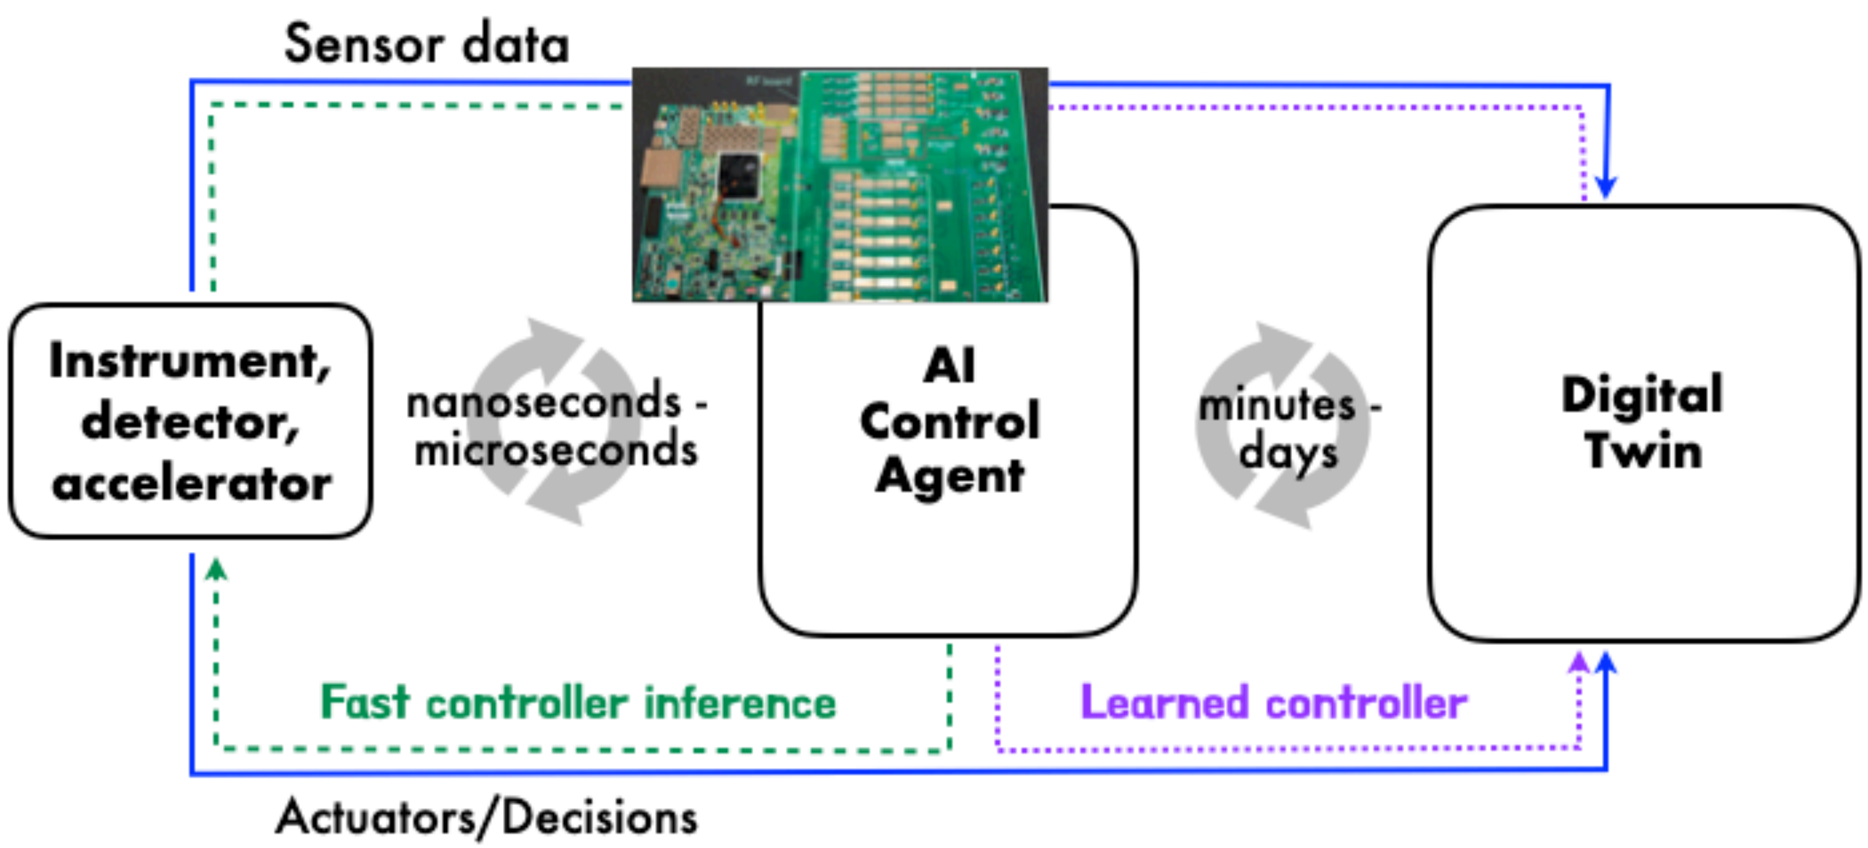

N. Tran

Summary

- A differentiable simulator for muon cooling channels is being developed
- Stay tuned for future updates