Building a Trustable Foundation Model for Collider Physics

Lead PI: Jesse Thaler (MIT)
Co-PIs from National Labs: Kyle Lee (Argonne) & Stefan Höche (Fermilab)
Co-PIs from Academia: Jennifer Roloff (Brown) & Matt LeBlanc (Brown)

Focus Area 14.A: Foundation Models of Particle Interactions and Cosmic Physics (HEP, NP)

FYI: This slide deck was converted by Claude from a pre-proposal to Argonne, but it appears that competition at Argonne is even more intense than at MIT! We are also exploring Brown as a possible lead, and the possibility of submitting under 19.A or 19.C.

Project Overview

A gap is opening between experimental precision and theoretical control that threatens to become the defining limitation of collider physics. The HL-LHC will results with experimental uncertainties below the percent level, yet current Monte Carlo (MC) event generators lag behind state-of-the-art analytic QCD calculations. **Closing this trustability gap is the central challenge motivating this proposal.**

We propose to build a **foundation model for collider physics**: an AI-enabled framework, grounded in information theory, that absorbs all available QCD knowledge—fixed-order calculations, analytic resummation, nonperturbative modeling—into a single coherent, fully differential event-level prediction. The core idea, rooted in the maximum-entropy principle, is that every piece of theoretical information can be cast as a constraint: the optimization finds the unique prediction that maximizes uncertainty in what we don't know while exactly imposing what we do know.

The team has established viability in a published proof of concept (Assi, Höche, Lee, Thaler, PRL 135 (2025) 131901). Reweighted predictions improve not only for directly constrained observables but also for correlated observables never used in training—demonstrating genuine information transfer. The AI advantage: machine learning identifies the optimal constraint basis and navigates high-dimensional optimization landscapes across millions of simulated events. The framework delivers principled uncertainty quantification as a natural byproduct.

9-months work plan

For Genesis Phase I, we will develop the maximum-entropy reweighting concept into a clear AI workflow for discovery-scale collider theory:

- **(1) Ingest precision theory inputs**
  - High-accuracy perturbative calculations, analytic resummation alongside MC simulation data
- **(2) Learn optimal constraint representation**
  - Systematic observable basis selection and moment-family optimization
- **(3) Perform event-level reweighting**
  - Full uncertainty propagation from theoretical inputs into reweighted predictions
- **(4) Quantify AI advantage**
  - Scaling metrics measuring prediction accuracy, transfer to unconstrained observables, and robustness across prior variations with constraint set size, data volume, and compute

Key Phase I deliverable: Completing and publishing the systematic study of energy flow polynomials (EFPs) as an observable basis for precision constraints. Scope extension from e^+e^- benchmarks to hadron-collider processes for HL-LHC and EIC.

9-months deliverables & Metrics

- **1. Publication and software: Precision EFP reweighting framework**
  - Complete and publish systematic study of EFP-based constraint families
  - Demonstrate information saturation, transfer to untrained observables, robustness across prior variations
  - Develop modular reweighting pipeline integrating precision theory inputs, MC priors (Sherpa, Pythia, PanScales), ML-based optimization, and uncertainty propagation
- **2. AI advantage benchmarks and uncertainty-aware event samples**
  - Quantitative benchmarks: transfer from constrained to unconstrained observables, robustness, measurable gains over baseline generators
  - Scaling metrics: fidelity vs. constraint count, data volume, and compute
  - Produce reweighted event samples with fully correlated perturbative uncertainty bands for experimental use
- **3. Hadron-collider extension (stretch goal)**
  - Initial feasibility study for extending beyond e^+e^- to pp collisions using jet-based EFPs
  - Assessment of pathways to incorporate nonperturbative QCD constraints
  - Concrete roadmap for scaling to full HL-LHC and EIC physics programs (continuing into Phase II)

Why the Team Is Ready Now

This project builds directly on demonstrated results: a published PRL establishing the maximum-entropy reweighting framework, with promising preliminary results on EFP-based constraint learning already in hand. Working code, validated event samples (10^6 events in Sherpa 3.0), and established computational infrastructure are already in place.

Thaler (MIT/IAIFI): Leader in jet physics, EFPs, and AI for fundamental physics; directs NSF AI Institute (IAIFI); developed the foundation-model-for-collider-physics vision
Lee (Argonne): Expertise in QCD resummation, jet substructure, and energy correlators; access to Argonne computing resources
Höche (Fermilab): Principal developer of Sherpa event generator; expert in parton shower algorithms and precision matching
Roloff & LeBlanc (Brown): State-of-the-art QCD measurements on ATLAS and CMS; expertise in systematics, uncertainty quantification, and QCD parameter inference

Phase I funding supports dedicated effort from existing team members, **including graduate students and postdocs**. No new hires or additional equipment required to begin immediately.

Path to Phase II

A successful Phase I will establish the technical foundation for building the full foundation model for collider physics at discovery scale in Phase II:

- Extending to multi-process precision campaigns spanning HL-LHC and EIC physics programs
- Enabling precision analyses under realistic fiducial cuts with coherent uncertainty propagation
- Integrating nonperturbative QCD constraints from lattice calculations and hadronization models
- Deploying production-grade reweighted event samples for ATLAS, CMS, sPHENIX, future EIC detectors

The methodology—**using information theory to anchor flexible AI models to first-principles physics**—is broadly transferable to scientific domains from nuclear astrophysics to materials science, positioning the Genesis investment for cross-cutting impact.