Hey folks,

Here's what you might find interesting today.

Jet Tagging and Jet Substructure

[Search for dark matter in a signature with a four-prong large-radius jet in proton-proton collisions at sqrt(s) = 13 TeV]

This CMS analysis uses a **graph neural network**-based jet substructure tagger to identify four-prong large-radius jets from boosted dark-sector mediators, marking the first search for nonprompt DM candidates in a Lorentz-boosted topology. The GNN tagger is applied to reconstruct the decay of light DM mediators into four quarks within a single large-radius jet, alongside two stable DM particles producing missing transverse momentum. This might be interesting for you, Aritra, as it directly combines **jet tagging with novel ML architectures** and **jet substructure exploitation** in a real LHC search context. The background estimation is done entirely from data using control regions, which is a nice touch for model-agnostic methodology. No excess is observed, but the analysis demonstrates the viability of GNN-based substructure tagging for exotic boosted signatures.

[Benchmarking the Nearside Energy-Energy Correlators with Mellin Transform]

This paper tackles **jet substructure physics** from an analytical angle, using a Mellin-transform framework to describe nearside energy-energy correlators (EECs) at small angles across both perturbative and post-confinement regimes. A single transition scale parameter Lambda provides an excellent fit to ALEPH data across different energies at NNLO+NNLL accuracy. This might be interesting for you, Aritra, as it deepens the theoretical understanding of EECs -- a observable closely tied to jet substructure that ML methods frequently exploit. While there is no ML component here, the analytical benchmarks it provides could serve as a reference point for evaluating ML-based approaches to jet substructure, and the Mellin-transform framework offers a complementary perspective to data-driven jet substructure characterisation.

Novel ML Architectures for Detector Reconstruction

[Antineutron reconstruction in electromagnetic calorimeters with mixed-representation learning]

This paper introduces a **Mixed-representation Calorimetric Network (MrCAL)** that integrates visual and sequential representation branches within a unified object-detection architecture for antineutron reconstruction in ECALs. The approach improves momentum-direction precision by up to 96% and, for the first time, enables direct momentum-magnitude measurement from ECAL readouts alone, achieving roughly 17% resolution at 1 GeV/c. This might be interesting for you, Aritra, because it proposes a genuinely novel ML architecture that fuses two distinct representation modalities -- image-like and sequence-like -- motivated by the two distinct energy deposition patterns of penetrating antineutrons. The generalization tests across diverse physics processes and background environments are thorough. The design philosophy of physics-inspired representation choices is worth noting for anyone thinking about how to architect ML for heterogeneous detector signals.

[Machine Learning-Based Reconstruction for Resistive Silicon Sensors]

This work studies ML-based reconstruction for AC-LGAD sensors, using full-waveform information from correlated pads with LSTM-based models and topology-agnostic transformer architectures that incorporate pad coordinates. The transformer extension is designed to support arbitrary pad counts and geometries, which is a thoughtful architectural choice for future sensor designs. This is interesting for you, Aritra, as it demonstrates how **novel ML architectures** (transformers with coordinate inputs) can be applied to detector-level reconstruction problems where the geometry is non-trivial. You might also find this relevant, Jingjing, as the paper touches on **ML uncertainty quantification** considerations -- the models are tested for FPGA deployment and bandwidth reduction, which connects to practical workflow automation in detector readout chains. The preservation of approximately 10 micron position resolution for 500 micron pitched sensors is a concrete performance benchmark.

AI Verification and Autonomous ML in Physics

Are We Ready for AI-Driven Discovery? AI Verification Before the Next Fundamental Physics Breakthrough

This review from the VERaiPHY initiative synthesises frameworks for rigorous ML assessment across particle physics, astrophysics, and cosmology, focusing on when verification becomes essential as ML systems grow more autonomous. The paper emphasises fundamental limitations -- inductive bias is unavoidable, sample complexity bounds learning, and experimental constraints limit discovery -- and discusses physicists' evolving role as both experimental designers and evaluators. This is interesting for you, Jingjing, as it directly addresses **agentic AI and workflow automation** concerns: the paper contextualises ML within the statistical discovery workflow and asks what happens when ML systems operate with increasing autonomy in hypothesis testing and inference. The discussion of inductive bias and sample complexity is also relevant to **simulation-based inference** reliability, where these issues manifest concretely. You might also find this relevant, Aritra, particularly the framing around **anomaly detection** and model-agnostic searches, where verification of unsupervised methods is an open challenge. Vinicius Mikuni and Lukas Heinrich are co-authors, both well-known for ML-in-HEP work, which adds credibility to the systematic framework proposed here.

Until the next announcement!

P.S. In the most recent submission window: hep-ph: 51 papers hep-ex: 26 papers hep-th: 43 papers