

ML4Jets2024

Monday 4 November 2024 - Friday 8 November 2024

LPNHE, Paris, France

Book of Abstracts

Contents

(R) Lorentz Group Equivariant Autoencoders	1
The Phase Space Distance Between Collider Events	1
Constraining the Higgs Potential with Neural Simulation-based Inference for Di-Higgs Production	2
Streamlined jet tagging network assisted by jet prong structure	2
Integrating Energy Flow Networks with Jet Substructure Observables for Enhanced Jet Quenching Studies	3
Classifying importance regions in Monte Carlo simulations with machine learning	3
Fast Perfekt: Regression-based refinement of fast simulation	4
Frequentist Uncertainties on Density Ratios with Ensembles	4
The versatility of flow-based fast calorimeter surrogate models	5
A Library for ML-based Fast Calorimeter Shower Simulation at Future Collider Experiments and Beyond	5
Towards Detector Agnostic Fast Calorimetry Simulation	6
New Physics Searches with Graph-Based Anomaly Detection in High-Energy Collisions	6
Calo4pQVAE: A calorimeter surrogate for high energy particle-calorimeter interactions using Dwave's Zephyr topology	7
Fusing physics principles and machine learning: inferring dark matter densities of galaxies using stellar catalogs with incomplete kinematic information	8
Semi-Supervised Permutation Invariant Particle-Level Anomaly Detection	8
Data-driven hadronization models	9
Learning Symmetry-Independent Jet Representation via Jet-Based Joint Embedding Predictive Architecture	9
Sweeping the Dust Away: An Unbiased Map of the Milky Way's Dark Matter and Gravitational Potential with Unsupervised Machine Learning	10
OmniLearn: A Method to Simultaneously Facilitate All Jet Physics Tasks	10
Full Event Particle-Level Unfolding with Variable-Length Latent Variational Diffusion	11

BitHEP –Are 1-Bit Networks all we need?	11
A Novel Approach to Training Foundation Models for Jet-Related Tasks Without Vector Quantization	12
Learning powerful jet representations via self-supervision	12
Jet Charge Classifiers	13
Machine learning the likelihoods	13
Differentiable MadNIS-Lite	14
Measurement of Jet Track Functions with OmniFold-based Binning Corrections in ATLAS Run 2 Data	14
The Fast Calorimeter Challenge 2022: Final Evaluation & Lessons Learned	15
Fair Universe HiggsML Uncertainty Challenge	15
Calibrating ATLAS calorimeter signals using an uncertainty-aware precision network	16
A Continuous Calibration of the ATLAS Flavor-Tagging Classifiers via Optimal Transporta- tion Maps	17
Machine Learning the Top Mass	17
OmniJet-alpha and beyond: foundation model updates	18
Point-Clouds based Diffusion Model on Hadronic Showers	18
Jet tagging with Lorentz-Equivariant Geometric Algebra Transformers	19
KAN we improve on HEP classification tasks? Kolmogorov-Arnold Networks applied to an LHC physics example	19
Resonant Searches as Cut and Count Experiments	20
Transformer for Energy Calibration in the ATLAS Electromagnetic Calorimeter	20
Generative transformers for learning point-cloud simulations	21
Learning the Simplicity of Scattering Amplitudes	22
Calibrating Bayesian Generative Machine Learning for Bayesiamplication	22
Large-Scale Pretraining and Finetuning for Efficient Jet Classification in Particle Physics	23
Towards Universal Unfolding using Denoising Diffusion	23
SKATR – A Self-Supervised Summary Transformer for the Square Kilometre Array	24
Meta-Learning Quantum Jet Properties with Quantum Generative Models	24
Generation of Air Shower Images for Imaging Air Cherenkov Telescopes using Diffusion Models	25
Event Generation with Lorentz-Equivariant Geometric Algebra Transformers	25

TRANSIT your events into a new mass: Fast background interpolation for semi-supervised anomaly detection searches	26
Pretrained event classification model for collider experiments	27
Efficient Particle Tracking and Pileup Mitigation with State space model	27
Faster Than Fast: Pushing the Limits of Simulation with Generative Models in HEP	28
Bridging the Generative Unfolding Gap	28
SkyCURTAINS: Model agnostic search for Stellar Streams with Gaia data	29
OmniFoldHI: Advanced ML Unfolding for Heavy-Ion Data	30
How to Unfold Top Decays	30
Enhancing generalization in high energy physics using white-box adversarial attacks . .	31
Supervised CWoLa: train supervised classifiers without background simulation	31
Parnassus: An Automated Approach to Accurate, Precise, and Fast Detector Simulation and Reconstruction	32
The Landscape of Unfolding with Machine Learning	32
Denoising Graph Super-Resolution for Improved Collider Event Reconstruction	33
CaloDREAM – Detector Response Emulation via Attentive flow Matching	33
Taming perturbation theory in QCD with Normalizing Flows	34
Exploring phase space with Flow Matching	34
Efficient SMEFT fits with neural importance sampling	35
WOTAN: Weakly-supervised Optimal Transport Attention-based Noise Mitigation	35
Advanced techniques for SBI in collider physics	36
Accelerating Graph-based Tracking Tasks with Symbolic Regression	36
Higher Resolution and Angular Conditioning for Normalizing-Flow-based Generation of Calorimeter Showers	37
Hadronic Top Quark Polarimetry with ParticleNet	37
CaloClouds III: Ultra-Fast Geometry-Independent Highly-Granular Calorimeter Simulation	38
Fast simulation of backgrounds at LHCb - a generalised tool	39
(R) Application of generative models for full-detector, whole-event simulated event generation and jet background subtraction	39
Gamma-ray spectrometry of fission fragments : ML analysis of multi-dimensional spectra	40
Heavy-Flavour Jet Tagging at LHCb Using Graph Neural Networks	40

The good, the bad, and the Bayesian?	41
The Fundamental Limit of Jet Tagging	41
Graph Neural Network-Based Track Finding in the LHCb Vertex Detector	42
Parameter Estimation with Neural Simulation-Based Inference in ATLAS	43
Efficient Resonant Anomaly Detection	43
Refining CMS Fast Simulation with ML-based regression	43
Introducing Aspen Open Jets: a real-world ML-ready dataset for jet physics	44
Uncertainty Quantification and Anomaly Detection with Evidential Deep Learning	44
Generating particle-clouds with discrete features using Markov jump processes	45
Transforming Flavour Tagging on ATLAS	45
Model-agnostic search for dijet resonances with anomalous jet substructure in proton-proton collisions at $\sqrt{s} = 13$ TeV with the CMS detector	46
ParticleNet: Calibration of the jet energy scale using pT regression with partial Run3 data collected by the CMS experiment	46
UParT: A unified approach for jet-based object identification in CMS in Run 3	47
Performance of the CNN-based tau identification algorithm with Domain Adaptation using Adversarial Machine Learning for Run 2	47
Deep learning on jet modification in the presence of the QGP background	48
Transformer networks for constituent-based b-jet calibration with the ATLAS detector	49
Jet Finding as a Real-Time Object Detection Task	49
(R) Generative Neural Networks for Reconstructing Parton-Level Jet Showers after Hadronization	49
Particle flow and flavor tagging with DNN for Higgs factories	50
ML assisted Event Reconstruction in the CMS Phase-2 High Granularity Calorimeter Endcap	50
Synergizing Physics: Deep Learning Techniques for Time-of-Flight Reconstruction and Jet Tagging in High Energy Physics	51
Symbolic machine learning in physics	51
Experimental highlights: Edge AI for particle physics	51
Experimental highlights: The State of ML in LHC Science	52
MadNIS - A journey towards the first ML event generator	52
Implicit and Explicit Simulation-Based Inference for Cosmology	52

Theory Overview 53

Realtime reconstruction: Machine learning in reconstruction at LHC 53

Time Series Anomaly Detection: Overview and New Trends 54

Welcome 54

Closing remarks 54

Anomaly detection / 2**(R) Lorentz Group Equivariant Autoencoders****Authors:** Javier Mauricio Duarte¹; Nadya Chernyavskaya²; Raghav Kansal¹; Zichun Hao³¹ *Univ. of California San Diego (US)*² *CERN*³ *California Institute of Technology***Corresponding Authors:** zichun.hao@cern.ch, nadezda.chernyavskaya@cern.ch, javier.mauricio.duarte@cern.ch, raghav.kansal@cern.ch

There has been significant work recently in developing machine learning (ML) models in high energy physics (HEP) for tasks such as classification, simulation, and anomaly detection. Often these models are adapted from those designed for datasets in computer vision or natural language processing, which lack inductive biases suited to HEP data, such as equivariance to its inherent symmetries. Such biases have been shown to make models more performant and interpretable, and reduce the amount of training data needed. To that end, we develop the Lorentz group autoencoder (LGAE), an autoencoder model equivariant with respect to the proper, orthochronous Lorentz group $SO^+(3, 1)$, with a latent space living in the representations of the group. We present our architecture and several experimental results on jets at the LHC and find it outperforms graph and convolutional neural network baseline models on several compression, reconstruction, and anomaly detection metrics. We also demonstrate the advantage of such an equivariant model in analyzing the latent space of the autoencoder, which can improve the explainability of potential anomalies discovered by such ML models.

Track:

Anomaly detection

Mixed contributions / 3**The Phase Space Distance Between Collider Events****Author:** Tianji Cai¹**Co-authors:** Andrew Larkoski²; Giacomo Koszegi³; Junyi Cheng⁴; Nathaniel Craig¹ *SLAC National Accelerator Laboratory*² *UCLA*³ *UCSB*⁴ *Harvard University***Corresponding Authors:** larkoski@g.ucla.edu, koszegi@physics.ucsb.edu, tianjic@stanford.edu, ncraig@physics.ucsb.edu, junyi_cheng@g.harvard.edu

How can one fully harness the power of physics encoded in relativistic N -body phase space? Topologically, phase space is isomorphic to the product space of a simplex and a hypersphere and can be equipped with explicit coordinates and a Riemannian metric. This natural structure that scaffolds the space on which all collider physics events live opens up new directions for machine learning applications and implementation. Here we present a detailed construction of the phase space manifold and its differential line element, identifying particle ordering prescriptions that ensure that the metric satisfies necessary properties. We apply the phase space metric to several binary classification tasks, including discrimination of high-multiplicity resonance decays or boosted hadronic decays of electroweak bosons from QCD processes, and demonstrate powerful performance on simulated data. Our work demonstrates the many benefits of promoting phase space from merely a background on which calculations take place to being geometrically entwined with a theory's dynamics.

Track:

Theory

Unfolding & Inference / 4

Constraining the Higgs Potential with Neural Simulation-based Inference for Di-Higgs Production

Authors: Benjamin Nachman¹; Radha Mastandrea²; Tilman Plehn^{None}

¹ *LBNL*

² *University of California, Berkeley*

Corresponding Authors: plehn@uni-heidelberg.de, rmastand@berkeley.edu, benjamin.philip.nachman@cern.ch

Determining the form of the Higgs potential is one of the most exciting challenges of modern particle physics. Higgs pair production directly probes the Higgs self-coupling and should be observed in the near future at the High-Luminosity LHC. We explore how to improve the sensitivity to physics beyond the Standard Model through per-event kinematics for di-Higgs events. In particular, we employ machine learning through simulation-based inference to estimate per-event likelihood ratios and gauge potential sensitivity gains from including this kinematic information. In terms of the Standard Model Effective Field Theory, we find that adding a limited number of observables can help to remove degeneracies in Wilson coefficient likelihoods and significantly improve the experimental sensitivity.

Track:

Anomaly detection

Tagging / 6

Streamlined jet tagging network assisted by jet prong structure

Authors: Ahmed Hamed Ali Hammad¹; Mihoko Nojiri²

¹ *KEK*

² *Theory Center, IPNS, KEK*

Corresponding Authors: nojiri@post.kek.jp, ahhammad@cern.ch

Attention-based transformer models have become increasingly prevalent in collider analysis, offering enhanced performance for tasks such as jet tagging. However, they are computationally intensive and require substantial data for training. In this paper, we introduce a new jet classification network using an MLP mixer, where two subsequent MLP operations serve to transform particle and feature tokens over the jet constituents. The transformed particles are combined with subjet information using multi-head cross-attention so that the network is invariant under the permutation of the jet constituents. We utilize two clustering algorithms to identify subjects: the standard sequential re-combination algorithms with fixed radius parameters and a new IRC-safe, density-based algorithm of dynamic radii based on HDBSCAN. The proposed network demonstrates comparable classification performance to state-of-the-art models while boosting computational efficiency drastically. Finally,

we evaluate the network performance using various interpretable methods, including centred kernel alignment and attention maps, to highlight network efficacy in collider analysis tasks.

Track:

Tagging (Classification)

Tagging / 8

Integrating Energy Flow Networks with Jet Substructure Observables for Enhanced Jet Quenching Studies

Authors: Guilherme Milhano¹; João A. Gonçalves²

¹ *LIP-Lisbon & CERN TH*

² *LIP - IST*

Corresponding Authors: guilherme.milhano@cern.ch, jgoncalves@lip.pt

The phenomena of Jet Quenching, a key signature of the Quark-Gluon Plasma (QGP) formed in Heavy-Ion (HI) collisions, provides a window of insight into the properties of this primordial liquid. In this study, we rigorously evaluate the discriminating power of Energy Flow Networks (EFNs), enhanced with substructure observables, in distinguishing between jets stemming from proton-proton (pp) and jets stemming from HI collisions. This work is yet another step towards separating significantly quenched jets from relatively unmodified ones on a per-jet basis, which would enable increasingly more precise measurements of QGP properties. We have analyzed simple Energy Flow Networks (EFNs) and subsequently augmented them with global features such as N-Subjettiness observables and Energy Flow Polynomials (EFPs). Our primary objective is to gauge the power of these approaches in the context of Jet Quenching. Initial evaluations using Linear Discriminant Analysis (LDA) set a performance baseline, which is further enhanced through simple Deep Neural Networks (DNNs), capable of capturing non-linear relations in the data. Integrating EFPs and N-Subjettiness observables into EFNs results in the most performant model over this task, achieving state-of-the-art ROC AUC values of approximately 0.84, a very considerable value given that both medium response and underlying event contamination effects are taken into account.

Track:

Tagging (Classification)

Event generation / 9

Classifying importance regions in Monte Carlo simulations with machine learning

Authors: Kayoung Ban¹; Myeonghun Park²; Raymundo Ramos¹

¹ *Korea Institute for Advanced Study*

² *Seoultech*

Corresponding Author: raalraan@gmail.com

We attempt to extend the typical stratification of parameter space used during Monte Carlo simulations by considering regions of arbitrary shape. Such regions are defined by directly using their importance for the simulation, for example, a likelihood or scattering amplitude. In particular, we

consider the possibility that the parameter space may be high dimensional and the simulation costly to compute. With this in mind, we suggest using data already obtained from the simulation to train a neural network to separate a larger set of points into guessed regions. The simulation would later be applied only on points that are deemed important for the final result, for example, variance reduction. We will discuss the particularities and complications of dividing the parameter space in this way and the role of the neural network in this process. Moreover, we illustrate the process with a few examples, including scattering and event generation, and compare with other known techniques for Monte Carlo simulations.

Track:

Detector simulation & event generation

Detector Simulation / 10

Fast Perfekt: Regression-based refinement of fast simulation

Authors: Lars Stietz¹; Moritz Jonas Wolf²

Co-authors: Patrick Louis S Connor³; Peter Schleper²; Samuel Louis Bein²

¹ *Hamburg University of Technology (DE)*

² *Hamburg University (DE)*

³ *University Hamburg (DE)*

Corresponding Authors: patrick.connor@desy.de, lars.stietz@tuhh.de, samuel.bein@cern.ch, peter.schleper@physik.uni-hamburg.de, moritz.wolf@cern.ch

As data sets grow in size and complexity, simulated data play an increasingly important role in analysis. In many fields, two or more distinct simulation software applications are developed that trade off with each other in terms of accuracy and speed. The quality of insights extracted from the data stand to increase if the accuracy of faster, more economical simulation could be improved to parity or near parity with more resource-intensive but accurate simulation. We present Fast Perfekt, a machine-learned regression-based model for refining fast simulations that employs residual neural networks. A deterministic network is trained using a unique schedule that combines ensemble-based and pair-based loss functions. We explore this methodology in the context of an abstract analytical model and in terms of a realistic particle physics application based on jet properties in hadron collisions at the Large Hadron Collider.

Track:

Detector simulation & event generation

Uncertainties & Interpretability / 11

Frequentist Uncertainties on Density Ratios with Ensembles

Author: Sean Benevedes¹

¹ *Massachusetts Institute of Technology*

Corresponding Author: seanbenevedes@gmail.com

We propose a novel framework to obtain asymptotic frequentist uncertainties on machine learned classifier outputs by using model ensembles. With the well-known likelihood trick, this framework

can then be applied to the task of density ratio estimation to obtain statistically rigorous frequentist uncertainties on estimated likelihood ratios. As a toy example, we demonstrate that the framework can recover known likelihood ratios for simple Gaussian distributions, and that the resulting estimates and uncertainties for the likelihood ratios satisfy the desired coverage properties. We then apply this framework in a collider physics context, estimating the likelihood ratio between generated quark and gluon jets. Finally, we examine the use of the learned likelihood ratio and uncertainties for downstream statistical inference.

Track:

Uncertainties

Anomaly detection / 12

The versatility of flow-based fast calorimeter surrogate models

Authors: Ben Nachman¹; Claudius Krause²; David Shih^{None}; Haoxing Du^{None}; Ian Pang^{None}; Vinicius Massami Mikuni¹; Yunhao Zhu^{None}

¹ *Lawrence Berkeley National Lab. (US)*

² *HEPHY Vienna (ÖAW)*

Corresponding Authors: benjamin.philip.nachman@cern.ch, ian.pang@physics.rutgers.edu, vinicius.massami.mikuni@cern.ch, claudius.krause@oeaw.ac.at, dshih@physics.rutgers.edu

Normalizing flows have proven to be state-of-the-art for fast calorimeter simulation. With access to the likelihood, these flow-based fast calorimeter surrogate models can be used for other tasks such as unsupervised anomaly detection (arXiv:2312.11618) and particle incident energy calibration (arXiv:2404.18992) without any additional training costs. Using CaloFlow as an example, we show that the unsupervised anomaly detector is sensitive to a wide range of signals, while the calibration approach is prior-independent and has access to per-shower resolution information.

Track:

Detector Simulation / 13

A Library for ML-based Fast Calorimeter Shower Simulation at Future Collider Experiments and Beyond

Authors: Peter McKeown¹; Thorsten Lars Henrik Buss²; Frank-Dieter Gaede³; Gregor Kasieczka⁴; Anatolii Korol⁵; Katja Kruger³; Thomas Madlener⁵; Piyush Raikwar¹; Anna Zaborowska¹

¹ *CERN*

² *Universität Hamburg*

³ *Deutsches Elektronen-Synchrotron (DE)*

⁴ *Hamburg University (DE)*

⁵ *Deutsches Elektronen-Synchrotron (DESY)*

Corresponding Authors: anna.zaborowska@cern.ch, piyush.raikwar@cern.ch, peter.mckeown@cern.ch, thorsten.buss@desy.de, frank-dieter.gaede@cern.ch, thomas.madlener@cern.ch, gregor.kasieczka@cern.ch, akorol98.17@gmail.com, katja.kruger@cern.ch

Experiments at current and future colliders rely fundamentally on precise detector simulation. While traditional simulation approaches based on Monte Carlo techniques provide a high degree of physics

fidelity, they place an enormous burden on the available computational resources. This is particularly true of particle showers created in the calorimeters, which have been a focus of fast simulation efforts. Approaches based on deep generative models have proved to be particularly promising options to provide significant reductions in computing times, while also being sufficiently accurate. While numerous generative models designed for this task have been studied in the literature, less attention has been given to interfacing these models with the existing software ecosystems. This is an essential step if a model is to be eventually deployed in a production environment. It also provides a means to evaluate the physics performance of a fast shower simulation model after reconstruction, which ultimately dictates its suitability as a fast simulation tool.

In this contribution we describe DDFastShowerML, a library now available in Key4hep. This generic library provides a means of combining inference of generative models trained to simulate calorimeter showers with the DD4hep toolkit, using the fast simulation hooks that exist in Geant4. This makes it possible to simulate showers in realistically detailed detector geometries, such as those proposed for use at future colliders and for community challenges, while seamlessly combining full and fast simulation. The flexibility of the library will be demonstrated through examples of different models that have been integrated, and different detector geometries that have been studied. An overview of future plans will also be presented.

Track:

Detector simulation & event generation

Detector Simulation / 17

Towards Detector Agnostic Fast Calorimetry Simulation

Authors: Anna Zaborowska¹; Kyongmin Yeo²; Mikolaj Piorczynski^{None}; Peter McKeown¹; Piyush Raikwar¹; Renato Paulo Da Costa Cardoso¹

¹ *CERN*

² *IBM Research*

Corresponding Authors: peter.mckeown@cern.ch, piyush.raikwar@cern.ch, kyeo@us.ibm.com, mikolaj.piorczynski@cern.ch, anna.zaborowska@cern.ch, renato.cardoso@cern.ch

Calorimeter simulations based on Monte Carlo methods (Geant4), while accurate, are computationally expensive and time-consuming. In this regard, numerous efforts aim to accelerate these simulations faster via generative machine learning. Although these machine learning models tend to be faster than Geant4, their design demands a significant amount of time, computational resources, and manpower. These factors limit the use of such models for new detector geometries. To mitigate this issue, inspired by foundation models (GPT-3, Dall.E 2, OpenAI Whisper), we investigate the idea of reusing the knowledge acquired by our transformer-based diffusion model when trained on various detector geometries. Our model shows robust generalization to new detector geometries while requiring substantially less training time and data. Furthermore, we present our findings on applying various methods to address the well-known issue of slow sampling speed of diffusion models.

Track:

Detector simulation & event generation

Anomaly detection / 19

New Physics Searches with Graph-Based Anomaly Detection in High-Energy Collisions

Authors: Ines Isabel Gouveia Cipriano Piedade Moreira¹; Nuno Castro²; Rute Pedro¹

¹ *Laboratory of Instrumentation and Experimental Particle Physics (PT)*

² *LIP and University of Minho*

Corresponding Authors: rute.pedro@cern.ch, nuno.castro@cern.ch, ines.p.moreira@tecnico.ulisboa.pt

In the realm of high-energy physics, the use of graph network-based implementations offers the advantage of handling input datasets more closely aligned with their collection process in collider experiments. GNN-based approaches address the graph anomaly detection problem by utilizing information about graph features and structures to effectively learn to score anomalies. We represent a single jet as a graph, with each node corresponding to a hadronic constituent clustered into the jet. This approach enables the identification of anomalous jets, contributing to the detection of anomalies at the jet level.

We use simulated datasets of Dark Jets events as the benchmark signal model, where a heavy vector boson Z' mediator connects a Standard Model (SM) quark pair with a pair of dark quarks. These dark quarks shower and hadronize, generating dark jets. For the background, we consider QCD dijet events.

Our goal is to extract a vector embedding that maps high-dimensional graph information into a low-dimensional vector using convolution and pooling mechanisms. These mechanisms efficiently propagate and aggregate information across the graph. The resulting vector embedding serves as input to an AD method, such as a one-class Deep Support Vector Data Description (DeepSVDD) and Autoencoders, allowing for the prediction and classification of jets based on their anomaly scores. We compare the performance of these models with baseline deep learning approaches.

Track:

Anomaly detection

Detector Simulation / 20

Calo4pQVAE: A calorimeter surrogate for high energy particle-calorimeter interactions using Dwave's Zephyr topology

Authors: Abhishek Abhishek^{None}; Colin Warren Gay¹; Deniz Sogutlu²; Eric Paquet³; Geoffrey Fox⁴; Hao Jia¹; Ian Lu²; J. Quetzalcoatl Toledo-Marin²; Maximilian J Swiatlowski⁵; Roger Melko^{None}; Sebastian Gonzalez²; Sehmimul Hoque⁶; Wojtek Fedorko²

¹ *University of British Columbia (CA)*

² *TRIUMF*

³ *Digital Technologies Research Centre, National Research Council*

⁴ *University of Virginia*

⁵ *TRIUMF (CA)*

⁶ *University of Waterloo*

Corresponding Authors: eric.paquet@nrc-cnrc.gc.ca, rgmelko.uwaterloo@gmail.com, wfedorko@triumf.ca, maximilian.j.swiatlowski@cern.ch, akajal@triumf.ca, gcfexchange@gmail.com, s4hoque@uwaterloo.ca, j.toledo.mx@gmail.com, dsogutlu@triumf.ca, haojia@phas.ubc.ca, i.lu@mail.utoronto.ca, cgay@physics.ubc.ca, bastigonalez2000@gmail.com

One potential roadblock towards the HL-LHC experiment, scheduled to begin in 2029, is the computational demand of traditional collision simulations. Projections suggest current methods will require millions of CPU-years annually, far exceeding existing computational capabilities. Replacing the event showers module in calorimeters with quantum-assisted deep learning surrogates can help bridge the gap. We propose a quantum-assisted deep generative model that combines a variational autoencoder (VAE) with a Restricted Boltzmann Machine (RBM) embedded in its latent space. The RBM in latent space provides further expressiveness to the model. We leverage D-Wave's Zephyr Quantum Annealer as a quantum version of an RBM. Our framework sets a path towards utilizing

large-scale quantum simulations as priors in deep generative models and for high energy physics, in particular, to generate high-quality synthetic data for the HL-LHC experiments.

Track:

Detector simulation & event generation

Astro & Cosmo / 21

Fusing physics principles and machine learning: inferring dark matter densities of galaxies using stellar catalogs with incomplete kinematic information

Author: Sung Hak Lim¹

¹ *Rutgers University*

Corresponding Author: sunghak.lim@rutgers.edu

Galactic dynamics studies often face the challenge of incomplete kinematic information in stellar catalogs.

This incompleteness poses a significant challenge to a complete and model-independent measurement of local galactic dark matter densities using stellar dynamics.

This talk presents two innovative approaches that fuse physics principles with machine learning techniques, specifically normalizing flows for stellar phase space density estimation, to overcome these limitations.

First, we demonstrate a method for measuring dark matter density on the disk of the Milky Way by leveraging equilibrium assumptions to compensate for missing stars obscured by intergalactic dust clouds and estimate the selection function: the probability of stars being included in the catalog.

Second, we introduce a technique for measuring dark matter density in distant dwarf spheroidal galaxies, utilizing spherical symmetry and equivariant normalizing flows to infer missing distance and proper motion data.

By augmenting incomplete data with physically motivated constraints and sophisticated machine learning models, our methods enable comprehensive analyses of galactic dark matter distributions. We anticipate that these modern machine learning-based approaches will allow us to fully utilize the potential of current and future astronomical catalogs, significantly improving our understanding of galactic dark matter.

Track:

Astrophysics

Anomaly detection / 22

Semi-Supervised Permutation Invariant Particle-Level Anomaly Detection

Authors: Elena Busch¹; Gabriel Pinheiro Matos¹; Julia Lynne Gonski²; Ki Ryeong Park¹

¹ *Columbia University (US)*

² *SLAC National Accelerator Laboratory (US)*

Corresponding Authors: julia.gonski@cern.ch, ki.ryeong.park@cern.ch, gabriel.pinheiro.matos@cern.ch, elb2180@columbia.edu

The development of analysis methods that can distinguish potential beyond the Standard Model phenomena in a model-agnostic way can significantly enhance the discovery reach in collider experiments. However, the typical machine learning (ML) algorithms employed for this task require fixed length and ordered inputs that break the natural permutation invariance in collider events. To address this limitation, we have designed a semi-supervised anomaly detection tool that takes a variable number of particle-level inputs and leverages a signal model to encode this information into a permutation invariant, event-level representation via supervised training with a Particle Flow Network (PFN). We then utilize this encoding as input to an autoencoder to perform unsupervised ANomaly deTEction on particLe fLOW latent sPacE (ANTELOPE), classifying anomalous events based on a low-level and permutation invariant input modeling. In this talk, the ANTELOPE architecture will be presented, and its performance will be demonstrated on the LHC Olympics dataset. Future outlook and evolutions of the tool will be discussed.

Track:

Anomaly detection

Event generation / 23

Data-driven hadronization models

Co-authors: Ahmed Youssef ¹; Andrzej Konrad Siodmok ²; Ben Nachman ³; Christian Bierlich ⁴; Jure Zupan ¹; Manuel Szwec ; Michael Kent Wilkinson ⁵; Philip Ilten ⁵; Stephen Mrenna ; Tony Menzo

¹ *University of Cincinnati*

² *Jagiellonian University (PL)*

³ *Lawrence Berkeley National Lab. (US)*

⁴ *Lund University (SE)*

⁵ *University of Cincinnati (US)*

Corresponding Authors: christian.bierlich@hep.lu.se, philten@cern.ch, a.siodmok@cern.ch, menzoad@mail.uc.edu, benjamin.philip.nachman@cern.ch, szwecml@ucmail.uc.edu, mrenna@fnal.gov, michael.k.wilkinson@cern.ch, youssead@ucmail.uc.edu, jure.zupan@cern.ch

I'll discuss recent and ongoing developments related to the tuning and construction of machine-learning-based models of hadronization. Specifically, I will discuss efforts related to the extraction of microscopic hadronization dynamics from macroscopic 'jet-level' observables as well as efforts related to fully differentiable hadronization tunes utilizing post-hoc reweighting.

Based on 2203.04983, 2308.13459, 2311.09296 and ongoing work.

Track:

Detector simulation & event generation

Tagging / 26

Learning Symmetry-Independent Jet Representation via Jet-Based Joint Embedding Predictive Architecture

Co-authors: Farouk Mokhtar ¹; Haoyang Li ¹; Javier Mauricio Duarte ¹; Subash Katel ²; Zihan Zhao ¹

¹ *Univ. of California San Diego (US)*

² *UCSD*

Corresponding Authors: hal113@ucsd.edu, skatel@ucsd.edu, z.zhao@cern.ch, javier.mauricio.duarte@cern.ch, fmokhtar@ucsd.edu

This study introduces an approach to learning augmentation-independent jet representations using a Jet-based Joint Embedding Predictive Architecture (J-JEPA). This approach aims to predict various physical targets from an informative context, using target positions as joint information. We study several methods for defining the targets and context, including grouping subjects within a jet, and grouping jets within a full collision event. As an augmentation-free method, J-JEPA avoids introducing biases that could harm downstream tasks, which often require invariance under augmentations different from those used in pretraining. This augmentation-independent training enables versatile applications, offering a pathway toward a cross-task foundation model. J-JEPA has the potential to excel in various jet-based tasks such as jet classification, energy calibration, and anomaly detection. Moreover, as a self-supervised learning algorithm, J-JEPA pretraining does not require labeled datasets, which can be crucial with the impending dramatic increase in computational cost for HL-LHC simulation. The reduced dependency of J-JEPA on extensive labeled data allows learning physically rich representations from unlabeled data and fine-tuning the downstream models with only a small set of labeled samples. In a nutshell, J-JEPA provides a less biased, cost-effective, and efficient solution for learning jet representations.

Track:

Tagging (Classification)

Astro & Cosmo / 27

Sweeping the Dust Away: An Unbiased Map of the Milky Way's Dark Matter and Gravitational Potential with Unsupervised Machine Learning

Author: Eric Putney¹

Co-authors: David Shih ; Matthew Buckley ; Sung Hak Lim ²

¹ *Rutgers, The State University of New Jersey*

² *Rutgers University*

Corresponding Authors: mbuckley@physics.rutgers.edu, eputney@physics.rutgers.edu, dshih@physics.rutgers.edu, sunghak.lim@rutgers.edu

The dynamics of stars in our galaxy encode crucial information about the Milky Way's dark matter halo. However, extinction from foreground dust can bias studies of stellar populations. By solving the equilibrium collisionless Boltzmann equation with novel machine learning techniques, we estimate the unbiased 6-dimensional phase space density of an equilibrated stellar population and the underlying gravitational potential. Utilizing a normalizing flow-based estimate for the phase space density of stars from the Gaia space observatory, we derive the local gravitational potential of the Milky Way and correct the stellar phase space density for dust extinction. Our data-driven estimates align with recent 3-dimensional dust maps and analytic models of the Milky Way's potential. This measurement will enhance our understanding of the detailed structure and substructure of the Milky Way's dark matter halo.

Track:

Astrophysics

Foundation models / 28

OmniLearn: A Method to Simultaneously Facilitate All Jet Physics Tasks

Authors: Ben Nachman¹; Vinicius Massami Mikuni¹

¹ *Lawrence Berkeley National Lab. (US)*

Corresponding Authors: vinicius.massami.mikuni@cern.ch, benjamin.philip.nachman@cern.ch

Machine learning has become an essential tool in jet physics. Due to their complex, high-dimensional nature, jets can be explored holistically by neural networks in ways that are not possible manually. However, innovations in all areas of jet physics are proceeding in parallel. We show that large machine learning models trained for a jet classification task can improve the accuracy, precision, or speed of all other jet physics tasks. This is demonstrated by training a large model on a particular multiclass classification task and then using the learned representation for a different classification task, for a dataset with a different (full) detector simulation, for jets from a different collision system (versus), for generative models, for likelihood ratio estimation, and for anomaly detection. Our OmniLearn approach is thus a foundation model and is made publicly available for use in any area where state-of-the-art precision is required for analyses involving jets and their substructure.

Track:

Unfolding & Inference / 29

Full Event Particle-Level Unfolding with Variable-Length Latent Variational Diffusion

Authors: Aishik Ghosh¹; Alexander Shmakov¹; Daniel Whiteson¹; Kevin Thomas Greif¹; Michael James Fenton¹

¹ *University of California Irvine (US)*

Corresponding Authors: michael.james.fenton@cern.ch, alexander.shmakov@cern.ch, kgreif@uci.edu, daniel.whiteson@cern.ch, aishik.ghosh@cern.ch

The measurements performed by particle physics experiments must account for the imperfect response of the detectors used to observe the interactions. One approach, unfolding, statistically adjusts the experimental data for detector effects. Recently, generative machine learning models have shown promise for performing unbinned unfolding in a high number of dimensions. However, all current generative approaches are limited to unfolding a fixed set of observables, making them unable to perform full-event unfolding in the variable dimensional environment of collider data. This talk presents a novel modification of the variational latent diffusion model (VLD) approach to generative unfolding, which allows for unfolding of high- and variable-dimensional feature spaces. The performance of this method is evaluated in the context of semi-leptonic top quark pair production at the Large Hadron Collider.

Track:

Unfolding

Detector Simulation / 30

BitHEP –Are 1-Bit Networks all we need?

Authors: Claudius Krause¹; Daohan Wang²; Ramon Winterhalder³

¹ *HEPHY Vienna (ÖAW)*

² *HEPHY ÖAW*

³ *Università degli Studi di Milano*

Corresponding Authors: rpw@mailbox.org, claudius.krause@oeaw.ac.at, daohan.wang@oeaw.ac.at

With the rise of modern and complex neural network architectures, there is a growing need for fast and memory-efficient implementations to avoid computational bottlenecks in high-energy physics (HEP). We explore the performance of the BITNET architecture in state-of-the-art HEP applications, focusing on classification, regression and generative modeling tasks. Specifically, we apply BITNET to the CaloINN for fast calorimeter shower simulations, MADNIS for neural importance sampling, P-DAT for quark/gluon discrimination, and SMEFTNet for decay plane angle regression. Additionally, we incorporate Bayesian networks to model the uncertainties in BITNET's predictions. Our results demonstrate that BITNET consistently achieves competitive performance across these diverse applications while reducing the computational resources.

Track:

Detector simulation & event generation

Foundation models / 31

A Novel Approach to Training Foundation Models for Jet-Related Tasks Without Vector Quantization

Author: Masahiro Morinaga¹

Co-authors: Junichi Tanaka¹; Masahiko Saito¹

¹ *University of Tokyo (JP)*

Corresponding Authors: masahiro.morinaga@cern.ch, masahiko.saito@cern.ch, junichi.tanaka@cern.ch

This study proposes a new method for training foundation models designed explicitly for jet-related tasks. Like those seen in large language models, a foundation model is a pre-trained model that can be fine-tuned for various applications and is not limited to a specific task. Previous approaches often involve randomly masking inputs, such as tracks within a jet, and then predicting the masked parts. However, unlike methods in other fields like image recognition and point clouds, these proposed techniques show less improvement in accuracy for downstream tasks as the amount of training data increases when compared to models trained from scratch.

Most existing methods heavily rely on vector quantization, which is crucial in determining accuracy. In High Energy Physics (HEP), input variables often have highly skewed distributions, making them poorly suited for vector quantization. Additionally, vector quantization using neural networks is known to be very unstable during training.

In response to these challenges, we propose a method that reconstructs masked inputs without using vector quantization. To reduce biases introduced by the model architecture, we use a LLaMA-type Transformer. This approach aims to evaluate the effectiveness of pre-training methods that do not rely on HEP-specific knowledge. We also discuss the results of pre-training and fine-tuning using the JetClass dataset.

Track:

Foundation models / 32

Learning powerful jet representations via self-supervision

Authors: Congqiao Li¹; Huilin Qu²; Qibin Liu³; Shudong Wang⁴

¹ *Peking University (CN)*

² *CERN*

³ *Tsung-Dao Lee Institute (CN) & Shanghai Jiao Tong University (CN)*

⁴ *Chinese Academy of Sciences (CN)*

Corresponding Authors: shudong.wang@cern.ch, congqiao.li@cern.ch, qibin.liu@cern.ch, huilin.qu@cern.ch

We propose a new approach to learning powerful jet representations directly from unlabelled data. The method employs a Particle Transformer to predict masked particle representations in a latent space, overcoming the need for discrete tokenization and enabling it to extend to arbitrary input features beyond the Lorentz four-vectors. We demonstrate the effectiveness and flexibility of this method in several downstream tasks, including jet tagging and anomaly detection. Our approach provides a new path to a foundation model for particle physics.

Track:

Tagging / 33

Jet Charge Classifiers

Author: Rabia Husain^{None}

Co-authors: Katherine Fraser¹; Noah McNeal¹; Rashmish Mishra¹

¹ *Harvard University*

Corresponding Authors: noahmcneal@g.harvard.edu, rashmishmishra@fas.harvard.edu, rhusain@g.harvard.edu, kfraser@berkeley.edu

While there has been tremendous progress on jet classification in the last decade, classifying samples which are very similar is still an open problem. One example of this is tagging up vs. down-quark initiated jets, which historically have utilized the observable p_T weighted jet charge directly or as an input to neural networks. In this work, we explore whether this trend persists when adding jet charge to classifiers with the state of the art performance on other samples. Specifically, we modify the inputs to LorentzNet, ParT, and PELICAN, which utilize GNNs and transformers. We find two major takeaways: particle level charge or particle ID information greatly improves classification and the results are insensitive to the specific p_T weight in particle level jet charge, unlike older ones.

Track:

Tagging (Classification)

Theorie / 34

Machine learning the likelihoods

Author: Rafal Maselek^{None}

Corresponding Author: r.maselek@uw.edu.pl

In recent years, the ATLAS collaboration has provided full statistical models for some of their analyses, enabling highly precise reinterpretation of experimental limits. These models account for multiple nuisance parameters and correlations between signal bins, but their complexity often leads to lengthy computation times. This project aims to develop a method for efficient yet accurate reinterpretation of experimental results in phenomenological studies. Specifically, we are training Deep Neural Networks (DNNs) to perform likelihood interpolation, serving as surrogates for full statistical models. This approach can reduce computation times by several orders of magnitude while maintaining high precision.

In my talk, I will introduce the project and present recent advancements, including the development of a framework for generating data with Markov Chain Monte Carlo (MCMC) methods, training Neural Networks to interpolate likelihoods, and validating these models on real-world analyses. Our approach has been tested on several experimental analyses, demonstrating promising results. The long-term goal is to create a publicly available and maintainable database of trained machine learning models that can be integrated into various reinterpretation tools, providing a valuable resource for the particle physics community.

Track:

Theory

Event generation / 35

Differentiable MadNIS-Lite

Authors: Olivier Mattelaer¹; Ramon Winterhalder²; Theo Heimel³; Tilman Plehn^{None}

¹ *UCLouvain*

² *Università degli Studi di Milano*

³ *Heidelberg University*

Corresponding Authors: heimel@thphys.uni-heidelberg.de, plehn@uni-heidelberg.de, olivier.mattelaer@uclouvain.be, rpw@mailbox.org

Differentiable programming opens exciting new avenues in particle physics, also affecting future event generators. These new techniques boost the performance of current and planned MadGraph implementations. Combining phase-space mappings with a set of very small learnable flow elements, MadNIS-Lite, can improve the sampling efficiency while being physically interpretable. This defines a third sampling strategy, complementing VEGAS and the full MadNIS.

Track:

Detector simulation & event generation

Unfolding & Inference / 36

Measurement of Jet Track Functions with OmniFold-based Binning Corrections in ATLAS Run 2 Data

Author: Jingjing Pan¹

¹ *Yale University (US)*

Corresponding Author: jingjing.pan@cern.ch

Measurements of jet substructure are key to probing the energy frontier at colliders, and many of them use track-based observables which take advantage of the angular precision of tracking detectors. Theoretical calculations of track-based observables require “track functions”, which characterize the transverse momentum fraction r_q carried by charged hadrons from a fragmenting quark or gluon. This work presents a direct measurement of r_q distributions in dijet events from the 140 fb^{-1} of $\sqrt{s} = 13 \text{ TeV}$ proton-proton collisions collected by the ATLAS detector. The data are corrected for detector effects using a machine learning-based method named OmniFold. The scale evolution of the moments of the r_q distribution provides direct access to non-linear renormalization group evolution equations of QCD, and is compared with analytic predictions. When incorporated into future theoretical calculations, these results will enable a precision program of theory-data comparison for track-based jet substructure observables.

Track:

Unfolding

Mixed contributions / 37

The Fast Calorimeter Challenge 2022: Final Evaluation & Lessons Learned

Author: Claudius Krause¹

Co-authors: Anna Zaborowska²; Ben Nachman³; Dalila Salamani; David Shih; Gregor Kasieczka⁴; Michele Fauci Giannelli⁵

¹ *HEPHY Vienna (ÖAW)*

² *CERN*

³ *Lawrence Berkeley National Lab. (US)*

⁴ *Hamburg University (DE)*

⁵ *Chalmers*

Corresponding Authors: dshih@physics.rutgers.edu, benjamin.philip.nachman@cern.ch, gregor.kasieczka@cern.ch, ad_salamani@esi.dz, anna.zaborowska@cern.ch, faucci@chalmers.se, claudius.krause@oeaw.ac.at

I report the final results of the Fast Calorimeter Challenge 2022: 23 collaborations submitted 59 samples across all 4 datasets. I will show how these rank regarding various metrics judging shower quality, generation time, and other properties. From these results, I present the current, state-of-the-art, Pareto fronts for using deep generative models on high-dimensional datasets in high-energy physics. These results will shape the future of fast simulation in the analysis chain at the experiments. In addition, this dataset allowed us to study the evaluation of deep generative models in general. I will show the correlation between different quality metrics, such as a binary or multiclass classifiers or FPD/KPD scores, and discuss what we can learn from this for the future.

Track:

Detector simulation & event generation

Mixed contributions / 38

Fair Universe HiggsML Uncertainty Challenge

Authors: Aishik Ghosh¹; Ben Nachman²; Christopher Harris³; David Rousseau⁴; Ihsan Ullah⁵; Isabelle Guyon^{None}; Jordan Dudley⁶; Paolo Calafiura²; Peter Nugent⁶; Ragansu Chakkappai⁴; Sascha Diefenbacher²; Shih-Chieh Hsu⁷; Steven Farrell⁸; Wahid Bhimji^{None}; Yuan-Tang Chou⁹; Yulei Zhang¹⁰

¹ *University of California Irvine (US)*

² *Lawrence Berkeley National Lab. (US)*

³ *Unknown*

⁴ *IJCLab-Orsay*

⁵ *Challearn*

⁶ *Berkeley*

⁷ *University of Washington Seattle (US)*

⁸ *Lawrence Berkeley National Laboratory*

⁹ *University of Washington (US)*

¹⁰ *University of Washington, Seattle*

Corresponding Authors: jdudley@berkeley.edu, benjamin.philip.nachman@cern.ch, rousseau@lal.in2p3.fr, schsu@uw.edu, sdiefenbacher@lbl.gov, ihsan2131@gmail.com, guyon@clopinet.com, penugent@lbl.gov, christopher.harris@cern.ch, yulei.zhang@cern.ch, pcalafiura@lbl.gov, wbhimji@lbl.gov, yuan-tang.chou@cern.ch, sfarrell@lbl.gov, aishik.ghosh@cern.ch, ragansu.chakkappai@cern.ch

The Fair Universe project is organising the HiggsML Uncertainty Challenge, which has been running from Sep 2024 to 14th March 2025. It is a NeurIPS 2024 competition.

This HEP and Machine Learning competition is the first to strongly emphasise uncertainties: mastering uncertainties in the input training dataset and outputting credible confidence intervals.

The context is the measurement of the Higgs to tau+ tau- cross section like in HiggsML challenge on Kaggle in 2014, from a dataset of the 4-momentum signal state. Participants should design an advanced analysis technique that can not only measure the signal strength but also provide a confidence interval, from which correct coverage will be evaluated automatically from pseudo-experiments.

The confidence interval should include statistical and systematic uncertainties (concerning detector calibration, background levels, etc...). It is expected that advanced analysis techniques that can control the impact of systematics will perform best, thereby pushing the field of uncertainty-aware AI techniques for HEP and beyond.

The challenge is hosted on Codabench (an evolution of the popular Codalab platform); the significant resources needed (to run the thousands of pseudo-experiments needed) are possible thanks to using NERSC infrastructure as a backend.

Fair Universe NeuIPS whitepaper
competition site

Track:

Tagging (Classification)

Uncertainties & Interpretability / 39

Calibrating ATLAS calorimeter signals using an uncertainty-aware precision network

Authors: Theo Heimel¹; Peter Loch²; Tilman Plehn^{None}; Jad Mathieu Sardain²; Philip Velie^{None}; Lorenz Vogel¹

¹ *Heidelberg University*

² *University of Arizona (US)*

Corresponding Authors: heimel@thphys.uni-heidelberg.de, pmv8ev@virginia.edu, vogel@thphys.uni-heidelberg.de, loch@physics.arizona.edu, plehn@uni-heidelberg.de, jad.sardain@cern.ch

ATLAS explores modern neural networks for a multi-dimensional calibration of its calorimeter signal defined by clusters of topologically connected cells (topo-clusters). The Bayesian neural network (BNN) approach yields a continuous and smooth calibration function, including uncertainties on the calibrated energy per topo-cluster. In this talk the performance of this BNN-derived calibration is compared to an earlier calibration network and standard table-lookup-based calibrations. The BNN uncertainties are confirmed using repulsive ensembles and validated through the pull distributions. First results indicate that unexpectedly large learned uncertainties can be linked to particular detector regions.

Track:

Uncertainties

Mixed contributions / 40

A Continuous Calibration of the ATLAS Flavor-Tagging Classifiers via Optimal Transportation Maps

Authors: ATLAS Collaboration^{None}; Chris Pollard¹

Co-author: Malte Algren²

¹ *University of Warwick (GB)*

² *Universite de Geneve (CH)*

Corresponding Authors: chris.pollard@cern.ch, cpollard@cern.ch, malte.algren@unige.ch

A calibration of the ATLAS flavor-tagging algorithms using a new calibration procedure based on optimal transportation maps is presented. Simultaneous, continuous corrections to the b -, c -, and light flavor classification probabilities from jet tagging algorithms in simulation are derived for b -jets using $t\bar{t} \rightarrow b\bar{b}e\mu\nu\nu$ events. After application of the derived calibration maps, closure between simulation and observation is achieved for jet flavor observables used in ATLAS analyses of the LHC collision data. This continuous calibration opens up new possibilities for the future use of jet flavor information in LHC analyses and furthermore serves as a guide for deriving high-dimensional corrections to simulation via transportation maps, an important development for a broad range of inference tasks.

Track:

Tagging (Classification)

Uncertainties & Interpretability / 41

Machine Learning the Top Mass

Author: Katherine Fraser¹

Co-authors: Arindam Bhattacharya ; Matthew Schwartz

¹ *Harvard University*

Corresponding Authors: arindamb@g.harvard.edu, schwartz@g.harvard.edu, kfraser@berkeley.edu

Energy correlators have recently shown potential to improve the precision on the top mass precision measurement. However, existing measurement strategies still only use part of the information in the EEE distribution and rely on arbitrary shape choices. In this talk, we explore the ability of Machine Learning to effectively optimize shape choice and reduce error on the top mass. Specifically, we utilize several simulation based inference approaches (both supervised and unsupervised) to learn the full 3D distribution in energy correlator space, and then use both regression and (energy-weighted) likelihood based approaches to extract the optimal value of the error on the Top Mass from this 3D distribution.

Track:

Uncertainties

Foundation models / 43

OmniJet-alpha and beyond: foundation model updates

Authors: Anna Maria Cecilia Hallin¹; Gregor Kasieczka²; Joschka Birk²

¹ *University of Hamburg*

² *Hamburg University (DE)*

Corresponding Authors: gregor.kasieczka@cern.ch, joschka.birk@cern.ch, anna.hallin@uni-hamburg.de

OmniJet-alpha is the first cross-task foundation model for particle physics, demonstrating transfer learning between an unsupervised problem (jet generation) and a classic supervised task (jet tagging). While OmniJet-alpha is still at a prototype stage, the successful development of foundation models for physics data would represent a major breakthrough, as they have the potential to enhance physics performance while simultaneously reducing the necessary training time and data significantly. This talk will give an overview of the model, and present new developments and results using additional training sources.

Track:

Detector Simulation / 44

Point-Clouds based Diffusion Model on Hadronic Showers

Authors: Anatolii Korol^{None}; Frank-Dieter Gaede¹; Gregor Kasieczka²; Martina Mozzanica²

Co-authors: Erik Buhmann²; Katja Kruger¹; Thorsten Lars Henrik Buss³; William Korcari²

¹ *Deutsches Elektronen-Synchrotron (DE)*

² *Hamburg University (DE)*

³ *Universität Hamburg*

Corresponding Authors: frank-dieter.gaede@cern.ch, gregor.kasieczka@cern.ch, anatolii.korol@desy.de, erik.buhmann@desy.de, thorsten.buss@desy.de, katja.kruger@cern.ch, william.korcari@cern.ch

Simulating showers of particles in highly-granular detectors is a key frontier in the application of machine learning to particle physics. Achieving high accuracy and speed with generative machine learning models can enable them to augment traditional simulations and alleviate a major computing constraint.

Recent developments have shown how diffusion based generative shower simulation approach that do not rely on a fixed structure, but instead generates geometry-independent point clouds, are very efficient. We present a novel attention mechanism based extension to the CaloClouds 2 architecture that was previously used for simulating electromagnetic showers in the highly granular electromagnetic calorimeter of ILD with high precision. This attention mechanism allows to generate complex hadronic showers from pions with more pronounced substructure in the electromagnetic and hadronic calorimeter together. This is the first time that ML methods are used to generate hadronic showers in highly granular imaging calorimeters.

Track:

Detector simulation & event generation

Tagging / 47

Jet tagging with Lorentz-Equivariant Geometric Algebra Transformers

Authors: Huilin Qu¹; Jesse Thaler²; Johann Brehmer³; Jonas Simon Spinner^{None}; Tilman Plehn^{None}; Víctor Bresó Pla⁴

¹ *CERN*

² *MIT*

³ *NYU*

⁴ *University of Heidelberg*

Corresponding Authors: breso@thphys.uni-heidelberg.de, j.spinner@thphys.uni-heidelberg.de, plehn@uni-heidelberg.de, jthaler@mit.edu, johann.brehmer@nyu.edu, huilin.qu@cern.ch

Extracting scientific understanding from particle-physics experiments requires solving diverse learning problems with high precision and good data efficiency. We present the Lorentz Geometric Algebra Transformer (L-GATr), a new multi-purpose architecture for high-energy physics. L-GATr represents high-energy data in a geometric algebra over four-dimensional space-time and is equivariant under Lorentz transformations. At the same time, the architecture is a Transformer, which makes it versatile and scalable to large systems. In this talk we will focus on the application of L-GATr to the task of jet classification. We find that L-GATr is able to either match or outperform other baselines in the top tagging, quark-gluon and JetClass benchmarks. In addition, we further improve the accuracy of L-GATr in the top-tagging task by pretraining the model on the JetClass dataset and then fine-tuning on top tagging data. This strategy boosts the classification performance of L-GATr across every metric, establishing a new state of the art result.

Track:

Tagging (Classification)

Uncertainties & Interpretability / 48

KAN we improve on HEP classification tasks? Kolmogorov-Arnold Networks applied to an LHC physics example

Authors: Florian Alexander Mausolf¹; Jan Lukas Spaeh¹; Johannes Erdmann²

¹ *Rheinisch Westfaelische Tech. Hoch. (DE)*

² RWTH Aachen University

Corresponding Authors: janlukas.spah@physik.rwth-aachen.de, johannes.erdmann@cern.ch, florian.mausolf@cern.ch

Recently, Kolmogorov-Arnold Networks (KANs) have been proposed as an alternative to multilayer perceptrons, suggesting advantages in performance and interpretability. In this talk, we present the first application of KANs in high-energy physics, focusing on a typical binary classification task involving high-level features.

We study KANs with different depths and widths and include a comparison to multilayer perceptrons in terms of performance and number of trainable parameters.

We find that the learned activation functions of a one-layer KAN resemble the log-likelihood ratios of the input features. In deeper KANs, the activations in the first KAN layer differ from those in the one-layer KAN, which indicates that the deeper KANs learn more complex representations of the data. For the chosen classification task, we do not find that KANs are more parameter efficient.

However, small KANs may offer advantages in terms of interpretability that come at the cost of only a moderate loss in performance.

Track:

Interpretability

Mixed contributions / 49

Resonant Searches as Cut and Count Experiments

Authors: Alexander Mück¹; David Shih^{None}; Gregor Kasieczka²; Marie Hein¹; Michael Kramer³; Ranit Das⁴; Thorben Finke^{None}

¹ RWTH Aachen University

² Hamburg University (DE)

³ Rheinisch Westfaelische Tech. Hoch. (DE)

⁴ Rutgers University

Corresponding Authors: finke@physik.rwth-aachen.de, mueck@physik.rwth-aachen.de, gregor.kasieczka@cern.ch, rd804@scarletmail.rutgers.edu, mkraemer@physik.rwth-aachen.de, marie.hein@rwth-aachen.de, dshih@physics.rutgers.edu

Weakly supervised anomaly detection has been shown to have great potential for improving traditional resonance searches. We demonstrate that weak supervision offers a unique opportunity to turn a resonance search into a simple cut-and-count experiment, where the potential problem of background sculpting in a traditional bump hunt is absent. Moreover, the cut-and-count setting allows working with large background rejection rates, where weakly supervised methods typically show their greatest significance improvement. Our method also provides a simple way to benchmark weakly supervised anomaly detection approaches in an end-to-end application. We quantify the performance of such a cut-and-count search using the CWoLa and Cathode approaches on the LHC Olympics R&D dataset.

Track:

Anomaly detection

Reconstruction / 50

Transformer for Energy Calibration in the ATLAS Electromagnetic Calorimeter

Authors: ATLAS Collaboration^{None}; Ryan Roberts¹

¹ *University of California Berkeley (US)*

Corresponding Authors: benjamin.ryan.roberts@cern.ch, atlas-perf-egamma-conveners@cern.ch

The ATLAS experiment reconstructs electrons and photons from clusters of energy deposits in the electromagnetic calorimeter. The reconstructed electron and photon energy must be corrected from the measured energy deposits in the clusters to account for energy loss in passive material upstream of the calorimeter, in the passive material in the calorimeter, out of cluster energies and leakage in the hadronic calorimeter. This correction is performed by a machine learning algorithm trained on Monte Carlo simulations to predict the true electron or photon energy based on a set of inputs describing the reconstructed electromagnetic cluster. This is complicated by the difficult pileup conditions observed in Run 3 and anticipated at the HL-LHC. It has been noted that a graph representation can naturally encode the irregular geometrical structure of calorimeter data. Transformer models are particularly suited for this task, as they can process graphs of arbitrary size and excel at capturing long-range dependencies between cells. Thanks to the self-attention mechanism, the transformer can weigh the importance of individual cells within the cluster and effectively model the complex relationships between them. We demonstrate that a transformer model has the potential to substantially improve the energy resolution with respect to the current binary decision tree (BDT) calibration, while adding resilience against pileup in nearly all kinematic regions studied. The transformer model is implemented in the SALT framework [1], commonly used on ATLAS for applications including jet flavor tagging, which allows a high degree of model customization for the calibration task and straightforward deployment to the ATLAS software framework.

[1] S. Van Stroud et al. 2024, <https://ftag-salt.docs.cern.ch/>

Track:

Reconstruction

Detector Simulation / 51

Generative transformers for learning point-cloud simulations

Authors: Anna Maria Cecilia Hallin¹; Gregor Kasieczka²; Henning Rose^{None}; Joschka Valentin Maria Birk^{None}; Martina Mozzanica²

¹ *University of Hamburg*

² *Hamburg University (DE)*

Corresponding Authors: martina.mozzanica@desy.de, joschka.birk@desy.de, henning.rose@desy.de, gregor.kasieczka@cern.ch, anna.hallin@uni-hamburg.de

We successfully demonstrate the use of a generative transformer for learning point-cloud simulations of electromagnetic showers in the International Large Detector (ILD) calorimeter. By reusing the architecture and workflow of the “OmniJet-alpha” model, this transformer predicts sequences of tokens that represent energy deposits within the calorimeter. This autoregressive approach enables the model to learn the sequence length of the point cloud, supporting a variable-length and realistic shower development. Furthermore, the tokenized representation allows the model to learn the shower geometry without being restricted to a fixed voxel grid.

Track:

Detector simulation & event generation

Theorie / 52

Learning the Simplicity of Scattering Amplitudes

Authors: Aurelien Dersy¹; Clifford Cheung^{None}; Matthew Schwartz^{None}

¹ *Harvard University*

Corresponding Authors: clifford.cheung@caltech.edu, adersy@g.harvard.edu, schwartz@g.harvard.edu

The simplification and reorganization of complex expressions lies at the core of scientific progress, particularly in theoretical high-energy physics. This work explores the application of machine learning to a particular facet of this challenge: the task of simplifying scattering amplitudes expressed in terms of spinor-helicity variables. We demonstrate that an encoder-decoder transformer architecture achieves impressive simplification capabilities for expressions composed of handfuls of terms. Lengthier expressions are implemented in an additional embedding network, trained using contrastive learning, which isolates subexpressions that are more likely to simplify. The resulting framework is capable of reducing expressions with hundreds of terms - a regular occurrence in quantum field theory calculations - to vastly simpler equivalent expressions. Starting from lengthy input expressions, our networks can generate the Parke-Taylor formula for five-point gluon scattering, as well as new compact expressions for five-point amplitudes involving scalars and gravitons.

Track:

Theory

Uncertainties & Interpretability / 53

Calibrating Bayesian Generative Machine Learning for Bayesian Amplification

Authors: Gregor Kasieczka¹; Mathias Trabs^{None}; Sascha Diefenbacher²; Sebastian Guido Bieringer³

¹ *Hamburg University (DE)*

² *Lawrence Berkeley National Lab. (US)*

³ *Hamburg University*

Corresponding Authors: sebastian.bieringer@desy.de, sdiefenbacher@lbl.gov, gregor.kasieczka@cern.ch, trabs@kit.edu

Generative models are on a fast track to becoming a mainstay in particle physics simulation chains, seeing active work towards adoption by nearly every large experiment and collaboration. However, the question of estimating the uncertainties and statistical expressiveness of samples produced by generative ML models is still far from settled.

Recently, combinations of generative and Bayesian machine learning have been introduced in particle physics for both fast detector simulation and inference tasks. These neural networks aim to quantify the uncertainty on the generated distribution originating from limited training statistics. The interpretation of a distribution-wide uncertainty however remains ill-defined. We show a clear scheme for quantifying the calibration of Bayesian generative machine learning models. For a Continuous Normalizing Flow applied to a low-dimensional toy example, we evaluate the calibration of Bayesian uncertainties from either a mean-field Gaussian weight posterior, or Monte Carlo sampling network weights, to gauge their behaviour on unsteady distribution edges. Well calibrated uncertainties can then be used to roughly estimate the number of uncorrelated truth samples that are equivalent to the generated sample and clearly indicate data amplification for smooth features of the distribution.

Track:

Uncertainties

Foundation models / 54

Large-Scale Pretraining and Finetuning for Efficient Jet Classification in Particle Physics

Author: Zihan Zhao¹

Co-authors: Farouk Mokhtar¹; Raghav Kansal¹; Haoyang Li¹; Javier Mauricio Duarte¹

¹ *Univ. of California San Diego (US)*

Corresponding Authors: raghav.kansal@cern.ch, hal113@ucsd.edu, javier.mauricio.duarte@cern.ch, fmokhtar@ucsd.edu, z.zhao@cern.ch

This study introduces an innovative approach to analyzing unlabeled data in high-energy physics (HEP) through the application of self-supervised learning (SSL).

Faced with the increasing computational cost of producing high-quality labeled simulation samples at the CERN LHC, we propose leveraging large volumes of unlabeled data to overcome the limitations of supervised learning methods, which heavily rely on detailed labeled simulations. By pretraining models on these vast, mostly untapped datasets, we aim to learn generic representations that can be finetuned with smaller quantities of labeled data. Our methodology employs contrastive learning with augmentations on jet datasets to teach the model to recognize common representations of jets, addressing the unique challenges of LHC physics.

Building on the groundwork laid by previous studies, our work demonstrates the critical ability of SSL to utilize large-scale unlabeled data effectively.

We showcase the scalability and effectiveness of our models by gradually increasing the size of the pretraining dataset and assessing the resultant performance enhancements.

Our results, obtained from experiments on two datasets—JetClass, representing unlabeled data, and Top Tagging, serving as labeled simulation data—show significant improvements in data efficiency, computational efficiency, and overall performance. These findings suggest that SSL can greatly enhance the adaptability of ML models to the HEP domain. This work opens new avenues for the use of unlabeled data in HEP and contributes to a better understanding of the potential of SSL for scientific discovery.

Track:

Tagging (Classification)

Unfolding & Inference / 55

Towards Universal Unfolding using Denoising Diffusion

Author: Pierre-Hugues Beauchemin¹

Co-authors: Camila Pazos¹; Martin Klassen¹; Shuchin Aeron²; Taritree Wongjirad²; Vincent Alexander Croft³

¹ *Tufts University (US)*

² *Tufts University*

³ *Nikhef National institute for subatomic physics (NL)*

Corresponding Authors: vincent.croft@cern.ch, c.pazos@cern.ch, martin.klassen@cern.ch, shuchin@eecs.tufts.edu, hugo.beauchemin@tufts.edu, taritree.wongjirad@tufts.edu

Correcting for detector effects in experimental data, particularly through unfolding, is critical for enabling precision measurements in high-energy physics. However, traditional unfolding methods face challenges in scalability, flexibility, and dependence on simulations. We introduce a novel approach to multidimensional particle-wise unfolding using conditional Denoising Diffusion Probabilistic Models (cDDPM). Our method utilizes the cDDPM for a non-iterative, flexible posterior sampling approach, incorporating distribution moments as conditioning information, which exhibits a strong inductive bias that allows it to generalize to unseen physics processes without explicitly assuming the underlying distribution. Our results highlight the potential of this method as a step towards a “universal” unfolding tool that reduces dependence on truth-level assumptions.

Track:

Unfolding

Astro & Cosmo / 56

SKATR – A Self-Supervised Summary Transformer for the Square Kilometre Array

Authors: Ayodele Ore^{None}; Caroline Heneka^{None}; Tilman Plehn^{None}

Corresponding Authors: ore@thphys.uni-heidelberg.de, heneka@thphys.uni-heidelberg.de, plehn@uni-heidelberg.de

The upcoming Square Kilometre Array (SKA) will bring about a new era of radio astronomy by allowing 3D imaging of the Universe during the periods of Cosmic Dawn and Reionisation. Machine learning promises to be a powerful tool to analyse the highly structured and complex signal, however accurate training datasets are expensive to simulate and supervised learning may not generalise. We introduce SKATR, a self-supervised vision transformer whose learned encoding can be cheaply adapted for downstream tasks on SKA maps. Focusing on regression and posterior inference of simulation parameters, we demonstrate that SKATR representations are near lossless. We also study how SKATR generalises to differently-simulated datasets and compare to fully-supervised baselines.

Track:

Astrophysics

Mixed contributions / 57

Meta-Learning Quantum Jet Properties with Quantum Generative Models

Authors: Cenk Tüysüz¹; Vincent Alexander Croft²; Yacine Haddad³

¹ DESY

² Nikhef National institute for subatomic physics (NL)

³ Northeastern University (US)

Corresponding Authors: yacine.haddad@cern.ch, ctuysuz@mail.desy.de, vincent.croft@cern.ch

Quantum Generative Models are emerging as a promising tool for modelling complex physical phenomena. In this work, we explore the application of Quantum Boltzmann Machines and Quantum

Generative Adversarial Networks to the intricate task of jet substructure modelling in high-energy physics. Specifically, we use these quantum frameworks to model the kinematics and corrections of the leading hadrons within a jet, focusing on accurately capturing quantum correlations. The aim is to evaluate whether quantum computing can reproduce the complex correlations observed in jets, which are challenging to simulate with classical methods. Our approach leverages a quantum-enhanced generative model to generate features that encapsulate the underlying quantum nature of jet evolution, incorporating quantum interference effects between hadrons. By studying this novel application of quantum generative models, we analyse their ability to outperform classical models in capturing such complex structures. We also investigate the impact of barren plateaus in training deep quantum circuits and propose strategies to mitigate their effects. Our empirical results provide insight into the potential of quantum computing in jet physics, paving the way for future applications in quantum-assisted substructure modelling.

Track:

Astro & Cosmo / 58

Generation of Air Shower Images for Imaging Air Cherenkov Telescopes using Diffusion Models

Author: Christian Elflein¹

Co-authors: Stefan Funk¹; Jonas Glombitza¹; Vinicius Massami Mikuni²; Ben Nachman²; Lark Wang²

¹ Erlangen Centre for Astroparticle Physics

² Lawrence Berkeley National Lab. (US)

Corresponding Authors: vinicius.massami.mikuni@cern.ch, s.funk@fau.de, benjamin.philip.nachman@cern.ch, christian.elflein@fau.de, jonas.glombitza@fau.de, lark_wang@berkeley.edu

The major goal of Imaging Atmospheric Cherenkov Telescopes (IACTs) is the investigation of gamma-ray sources through the detection of their induced air showers. For every detected gamma ray, there are up to 10000 cosmic ray protons present forming the background, which also needs to be studied. For a detailed understanding of the instrument for deriving its response to both gamma rays and protons, a significant number of simulations are required. These simulations are computationally extensive and time extensive, particularly for the simulation of proton-induced showers as their structure is of a more complex nature. Additionally, changes in the observation conditions also result in the need for new simulations. Thus, novel approaches to increase the efficiency and accelerate the shower simulations offer new prospects for astroparticle physics. Diffusion models have been established as the state of the art over the last years and demonstrated their effectiveness in fast event generation. In this work, we apply a score-based diffusion model to investigate the fast generation of IACT images using simulations of the High Energy Stereoscopic System (H.E.S.S.). The IACT camera features the FlashCam design, foreseen for the Cherenkov Telescope Array (CTA), with over 1500 pixels. The successful application of this machine learning model is verified through the analysis of several high- and low-level parameters that give information about the image and air shower properties. Furthermore, we compared the generated images of diffusion models to generative adversarial networks and found promising performance for the fast generation of IACT images.

Track:

Astrophysics

Event generation / 59

Event Generation with Lorentz-Equivariant Geometric Algebra Transformers

Authors: Huilin Qu¹; Jesse Thaler²; Johann Brehmer³; Jonas Simon Spinner^{None}; Pim De Haan³; Tilman Plehn^{None}; Víctor Bresó Pla⁴

¹ *CERN*

² *MIT*

³ *Qualcomm, CuspAI*

⁴ *University of Heidelberg*

Corresponding Authors: j.spinner@thphys.uni-heidelberg.de, huilin.qu@cern.ch, plehn@uni-heidelberg.de, breso@thphys.uni-heidelberg.de, jthaler@mit.edu

Extracting scientific understanding from particle-physics experiments requires solving diverse learning problems with high precision and good data efficiency. We propose the Lorentz Geometric Algebra Transformer (L-GATr), a new multi-purpose architecture for high-energy physics. L-GATr represents high-energy data in a geometric algebra over four-dimensional space-time and is equivariant under Lorentz transformations, the symmetry group of relativistic kinematics. At the same time, the architecture is a Transformer, which makes it versatile and scalable to large systems. We use L-GATr to construct the first Lorentz-equivariant generative network for LHC events. The continuous normalizing flow is trained with Riemannian flow matching, where we incorporate knowledge about challenging phase space features into the construction of the target velocity field. We discuss the role of symmetry breaking in the construction of the L-GATr generator. Across all performance metrics, the L-GATr generator surpasses equivariant and non-equivariant baselines, positioning it as a robust and innovative framework for pushing the boundaries of machine learning in particle physics.

Track:

Detector simulation & event generation

Anomaly detection / 60

TRANSIT your events into a new mass: Fast background interpolation for semi-supervised anomaly detection searches

Authors: Ivan Oleksiyuk¹; Svyatoslav Voloshynovskyy²; Tobias Golling³

¹ *UNIGE*

² *University of Geneva*

³ *Universite de Geneve (CH)*

Corresponding Authors: tobias.golling@unige.ch, ivan.oleksiyuk@gmail.com, svyatoslav.voloshynovskyy@unige.ch

We introduce TRANSIT, a conditional adversarial network for continuous interpolation of data. It is designed to construct a background data template for semi-supervised searches for new physics processes at the LHC, by smoothly transforming sideband events to match signal region mass distributions.

We demonstrate the performance of TRANSIT using the LHC Olympics R&D dataset. The method effectively captures non-linear mass correlations within given features and produces a template that offers competitive anomaly detection sensitivity compared to state-of-the-art (SotA) template generators. Additionally, the computational training time for TRANSIT is an order of magnitude lower than that of competing deep learning methods, making it particularly advantageous for analyses involving numerous signal regions and models.

Unlike most generative models, which must learn the full probability density distribution—i.e., the correlations between all variables—the proposed model only needs to learn a smooth conditional shift of the distribution. This simplifies the architecture and significantly enhances efficiency.

The absence of an informational bottleneck and the use of a residual architecture facilitate mass-uncorrelated features to pass through the network unchanged, while mass-correlated features are adjusted accordingly.

The proposed approach is based on a variational approximation of mutual information via adversarial decomposition, further contributing to its robustness and flexibility.

Track:

Anomaly detection

Tagging / 61

Pretrained event classification model for collider experiments

Authors: Haichen Wang¹; Joshua Anthony Ho¹; Shuo Han¹; Ryan Roberts²

¹ *Lawrence Berkeley National Lab. (US)*

² *University of California Berkeley (US)*

Corresponding Authors: benjamin.ryan.roberts@cern.ch, ho22joshua@berkeley.edu, shuo.han@cern.ch, haichen.wang@cern.ch

Analysis of collision data often involves training deep learning classifiers on very specific tasks and in regions of phase-space where the training datasets have limited statistics. Models pre-trained on a larger, more generic, sample may already have a useful representation of collider data which can be leveraged by many independent downstream analysis tasks. We introduce a class of pre-trained neural network models that can be fine-tuned for specific collider event classification tasks. These models are based on graph neural network architecture and have been trained on a large dataset of diverse simulated collision events for various classification and regression tasks. Our findings demonstrate that when fine-tuned for a new analysis task, the pre-trained model can outperform a classification model directly trained for that specific task. This improvement is particularly significant when the training sample for the downstream analysis task has limited statistics. In several tests, the pre-trained model also exhibits faster convergence during training, offering the potential to reduce overall time and energy consumption in scenarios that require repeated model training. Additionally, we present studies on the similarity of representations between the pre-trained model and models directly trained for the final analysis tasks.

Track:

Tagging (Classification)

Reconstruction / 62

Efficient Particle Tracking and Pileup Mitigation with State space model

Authors: Cheng Jiang¹; Huilin Qu²; Sitian Qian³; Yihui Lai⁴; Yongbin Feng⁵

¹ *The University of Edinburgh (GB)*

² *CERN*

³ *Peking University (CN)*

⁴ *University of Maryland (US)*

⁵ *Texas Tech University (US)*

Corresponding Authors: yongbin.feng@cern.ch, stqian@pku.edu.cn, chji4678@colorado.edu, yihui.lai@cern.ch, huilin.qu@cern.ch

Large-scale point cloud and long-sequence processing are crucial for high energy physics applications such as pileup mitigation and track reconstruction. The HL-LHC presents inevitable challenges to machine learning models, requiring both high stability and low computational complexity. Previous studies have primarily focused on graph-based approaches which are generally effective but often struggle with computational complexity. In this study, we introduce the state space model with several key improvements. For example, based on similar logic as the Kalman Filter, Mamba is used with customized depth-wise convolution and SSM blocks. Ideally, Mamba should have inference times as short as Gated MLP when sequences become longer. We also have integrated a new matrix mixer and local-sensitive architecture into Mamba to further improve the throughput while having the same performance. To better simulate future realistic scenarios, we emphasize the long sequences case, where many models suffer from high complexity. Preliminary results show better performance than previous graph and transformer approach on node-level classification and clear physics evaluation metrics improvement on most kinematics regions yet still achieving a much stronger error-complexity/performance-speed tradeoff.

Track:

Tagging (Classification)

Detector Simulation / 63

Faster Than Fast: Pushing the Limits of Simulation with Generative Models in HEP

Authors: Cheng Jiang¹; Huilin Qu²; Sitian Qian³

¹ *The University of Edinburgh (GB)*

² *CERN*

³ *Peking University (CN)*

Corresponding Authors: huilin.qu@cern.ch, stqian@pku.edu.cn, chji4678@colorado.edu

Fast event and detector simulation in high-energy physics using generative models provides a viable solution for generating sufficient statistics within a constrained computational budget, particularly in preparation for the High Luminosity LHC. However, many of these applications suffer from a quality/speed tradeoff. Diffusion models offer some of the best sampling quality but slow generation with large sampling steps. In our study, we replaced the traditional neural network backbone with a GBDT-based backbone to specifically address unstructured tabular data. This results in training and inference times for most high-level simulation tasks being sped up by orders of magnitude. The application can be extended to low-level feature simulation and conditioned generation with competitive performance. We also conducted a comprehensive scan of most mainstream samplers for standard score matching diffusion, achieving an $O(10)$ speedup with training-free methods. The new signal-to-noise ratio weighting and step-aware scheduler fine-tuning methods are introduced to enable most ODE samplers to perform well with around 10 evaluation steps.

Track:

Detector simulation & event generation

Unfolding & Inference / 64

Bridging the Generative Unfolding Gap

Authors: Anja Butter¹; Ben Nachman²; Nathan Huetsch³; Sascha Diefenbacher²; Sofia Palacios Schweitzer⁴; Tilman Plehn^{None}; Vinicius Massami Mikuni²

¹ *Centre National de la Recherche Scientifique (FR)*

² *Lawrence Berkeley National Lab. (US)*

³ *Heidelberg University, ITP Heidelberg*

⁴ *ITP, University Heidelberg*

Corresponding Authors: abutter@lpnhe.in2p3.fr, huetsch@thphys.uni-heidelberg.de, benjamin.philip.nachman@cern.ch, palacios@thphys.uni-heidelberg.de, vinicius.massami.mikuni@cern.ch, sdiefenbacher@lbl.gov, plehn@uni-heidelberg.de

Machine learning-based unfolding has started to establish itself as the go-to approach for precise, high-dimensional unfolding tasks. The current state-of-the-art unfolding methods can be divided into reweighting-based and generation-based methods. The latter of the two is comprised of conditional generative models, which generate new truth-level events from random noise conditioned on detector-level inputs, and of bridge-based models, which directly map events from detector- to truth-level.

Bridge-based models have always had the advantage of starting from a physically motivated distribution, rather than from random noise, placing their starting points innately closer to the desired results. However, the mappings learned by the bridges were often more akin to an optimal-transport mapping between detector-level and truth-level, rather than to the mapping prescribed by the detector.

We show recent developments in addressing this shortcoming and present a set of improved bridge models, which are able to learn the exact detector mapping, in the same way conditional generative models can, without sacrificing the inherent advantages of utilizing a physically motivated distribution. We demonstrate the efficacy of these new bridges on a synthetic example set and on a Z+jets dataset.

Track:

Unfolding

Astro & Cosmo / 65

SkyCURTAINS: Model agnostic search for Stellar Streams with Gaia data

Authors: David Shih^{None}; Debajyoti Sengupta¹; Johnny Raine^{None}; Stephen Brian Mulligan¹; Tobias Golling¹

¹ *Universite de Geneve (CH)*

Corresponding Authors: johnnyraine92@gmail.com, dshih@physics.rutgers.edu, stephen.mulligan@unige.ch, debajyoti.sengupta@unige.ch, tobias.golling@unige.ch

We introduce SkyCURTAINS, an adaptation of the CURTAINS method—a weakly supervised technique originally developed for anomaly detection in high-energy physics data—applied to data from the second Gaia Data Release (GDR2). SkyCURTAINS is employed to search for stellar streams, which appear as line-like overdensities against the background of the Milky Way. To validate the feasibility of this approach, we evaluate its performance on the recovery of the GD-1 stream, a well studied stellar stream for which truth labels are available. SkyCURTAINS achieves a purity of 75.4%, a 10% improvement over existing methods, while maintaining a signal efficiency of 37.9%. These results highlight the effectiveness of generic, data-driven, and model-agnostic approaches in addressing anomaly detection across distinct domains. Notably, due to the generic nature of the method, CURTAINS can detect various types of anomalies, including streams, globular clusters, and dwarf galaxies. The SkyCURTAINS method is only specialised by the final step in the algorithm, which applies a Hough transform to specifically search for line-like structures, which leaves open

the possibility of future searches for these other types of anomalies. The success of this study naturally suggests a follow up full sky scan that could potentially discover previously unknown stellar streams.

Track:

Astrophysics

Unfolding & Inference / 66

OmniFoldHI: Advanced ML Unfolding for Heavy-Ion Data

Authors: Adam Takacs¹; Alexandre Falcão²

¹ *Heidelberg University*

² *University of Bergen*

Corresponding Authors: takacs@thphys.uni-heidelberg.de, alexandre.falcao@uib.no

To compare collider experiments, measured data must be corrected for detector distortions through a process known as unfolding. As measurements become more sophisticated, the need for higher-dimensional unfolding increases, but traditional techniques have limitations. To address this, machine learning-based unfolding methods were recently introduced. In this work, we introduce OmniFoldHI, an improved version of the well-known algorithm [1], tailored for heavy-ion analyses. OmniFoldHI incorporates background counts, detector acceptances, efficiency, and uncertainties for real-analysis applications, and it works for an arbitrary number of observables. Besides removing detector effects, we demonstrate that unfolding can be used to subtract the high-multiplicity underlying event, which is crucial for jet-quenching analyses and phenomenology. With these enhancements, OmniFoldHI functions effectively even without additional background subtraction. To illustrate its capabilities, we apply OmniFoldHI to unfold up to a 7-dimensional jet-substructure observable, comparing it to traditional techniques and quantifying uncertainties. We present model-independent results, with training and testing performed using different event generators.

[1] Andreassen et. al, Phys. Rev. Lett. 124, 182001 (2020)

Track:

Unfolding

Unfolding & Inference / 67

How to Unfold Top Decays

Authors: Alexander Paasch¹; Dennis Schwarz²; Luigi Favaro^{None}; Roman Kogler³; Sofia Palacios Schweitzer⁴; Tilman Plehn^{None}

¹ *Hamburg University (DE)*

² *Austrian Academy of Sciences (AT)*

³ *DESY (DE)*

⁴ *ITP, University Heidelberg*

Corresponding Authors: plehn@uni-heidelberg.de, luigi.favaro2@gmail.com, roman.kogler@cern.ch, palacios@thphys.uni-heidelberg.de, dennis.schwarz@cern.ch, alexander.paasch@uni-hamburg.de

Many physics analyses at the LHC rely on algorithms to remove detector effect, commonly known as unfolding. Whereas classical methods only work with binned, one-dimensional data, Machine Learning promises to overcome both problems. Using a generative unfolding pipeline, we show how it can be build into an existing LHC analysis, designed to measure the top mass. We discuss the model-dependence of our algorithm, i.e. the bias of our measurement towards the top mass used in simulation and propose a method to reliably achieve unbiased results.

Track:

Unfolding

Tagging / 68

Enhancing generalization in high energy physics using white-box adversarial attacks

Author: Franck Rothen¹

Co-authors: Matthew Leigh²; Samuel Byrne Klein¹; Tobias Golling¹

¹ *Universite de Geneve (CH)*

² *University of Geneva*

Corresponding Authors: samuel.byrne.klein@cern.ch, matthew.leigh@cern.ch, tobias.golling@unige.ch, franck.rothen@unige.ch

Machine learning is becoming increasingly popular in the context of particle physics. Supervised learning, which uses labeled Monte Carlo simulations, remains one of the most widely used methods for discriminating signals beyond the Standard Model. However, this paper suggests that supervised models may depend excessively on artifacts and approximations from Monte Carlo simulations, potentially limiting their ability to generalize well to real data. This study aims to enhance the generalization properties of supervised models. It reviews the application of four distinct white-box adversarial attacks, in the context of classifying Higgs boson decay signals. The attacks are divided into two groups: weight space attacks, and feature space attacks. A dense network is used to compare these methods. To study and quantify the sharpness of the found local minima, this paper also presents two analysis methods: gradient ascent and reduced Hessian eigenvalue analysis. The results show that white-box adversarial attacks significantly improve generalization performance, though they also increase computational complexity.

Track:

Tagging (Classification)

Tagging / 69

Supervised CWoLa: train supervised classifiers without background simulation

Authors: Samuel Byrne Klein¹; Stephen Mulligan²; Tobias Golling¹

¹ *Universite de Geneve (CH)*

² *Universite de Genève*

Corresponding Authors: tobias.golling@unige.ch, samuel.byrne.klein@cern.ch, stephen.mulligan@unige.ch

Supervised deep learning methods have found great success in the field of high energy physics (HEP) and the trend within the field is to move away from high level reconstructed variables to low level detector features. However, supervised methods require labelled data, which is typically provided by a simulator. The simulations of HEP datasets become harder to validate and calibrate as we move to low level variables. In this work we show that the classification without labels paradigm can be used to enhance supervised searches for specific signal models by removing the need for background simulation when training supervised classifiers. When combined with a data driven background estimation technique this allows for dedicated searches for specific new physics processes to be performed using simulated signal only.

Track:

Tagging (Classification)

Detector Simulation / 70

Parnassus: An Automated Approach to Accurate, Precise, and Fast Detector Simulation and Reconstruction

Authors: Ben Nachman¹; Dmitrii Kobylanski²; Eilam Gross²; Etienne Dreyer²; Nathalie Soybelman²; Vinicius Massami Mikuni¹

¹ *Lawrence Berkeley National Lab. (US)*

² *Weizmann Institute of Science (IL)*

Corresponding Authors: nathalie.soybelman@cern.ch, dmitrii.kobylanski@cern.ch, etienne.dreyer@cern.ch, vinicius.massami.mikuni@cern.ch, benjamin.philip.nachman@cern.ch, eilam.work@gmail.com

Simulating particle physics data is an essential yet computationally intensive process in analyzing data from the LHC. Traditional fast simulation techniques often use a surrogate calorimeter model followed by a reconstruction algorithm to produce reconstructed objects. In this work, we introduce Particle-flow Neural Assisted Simulations (Parnassus), a deep learning-based method for generating these reconstructed objects. Our model takes as input a point cloud representing particles interacting with the detector and outputs a point cloud of reconstructed particles. By integrating detector simulation and reconstruction into a single step, we aim to reduce resource consumption and create fast surrogate models that can be applied both within and beyond large collaborations. We demonstrate this approach using a publicly available dataset of jets processed through the full simulation and reconstruction pipeline of the CMS experiment. Our results show that the model accurately replicates the CMS particle flow algorithm on the same events used for training and generalizes well to different jet momenta and types outside the training distribution.

Track:

Detector simulation & event generation

Mixed contributions / 71

The Landscape of Unfolding with Machine Learning

Authors: Alexander Shmakov¹; Anja Butter²; Ben Nachman³; Daniel Whiteson¹; Javier Mariño Villadamigo⁴; Kevin Thomas Greif¹; Michael James Fenton¹; Nathan Huetsch⁵; Sascha Diefenbacher³; Theo Heimel⁶; Tilman Plehn⁶; Vinicius Massami Mikuni³

¹ *University of California Irvine (US)*

² *Centre National de la Recherche Scientifique (FR)*

³ *Lawrence Berkeley National Lab. (US)*

⁴ *Institut für Theoretische Physik - University of Heidelberg*

⁵ *Heidelberg University, ITP Heidelberg*

⁶ *Heidelberg University*

Corresponding Authors: huetsch@thphys.uni-heidelberg.de, heimel@thphys.uni-heidelberg.de, abutter@lpnhe.in2p3.fr, daniel.whiteson@cern.ch, tilman.plehn@cern.ch, michael.james.fenton@cern.ch, sdiefenbacher@lbl.gov, alexander.shmakov@cern.ch, benjamin.philip.nachman@cern.ch, vinicius.massami.mikuni@cern.ch, marino@thphys.uni-heidelberg.de, kgreif@uci.edu

Recent innovations from machine learning allow for data unfolding, without binning and including correlations across many dimensions. We describe a set of known, upgraded, and new methods for ML-based unfolding. The performance of these approaches is evaluated on two benchmark datasets. We find that all techniques are capable of accurately reproducing the particle-level spectra across complex observables. Given that these approaches are conceptually diverse, they offer an exciting toolkit for a new class of measurements that can probe the Standard Model with an unprecedented level of detail and may enable sensitivity to new phenomena.

Track:

Unfolding

Reconstruction / 72

Denoising Graph Super-Resolution for Improved Collider Event Reconstruction

Author: Nilotpal Kakati¹

Co-authors: Eilam Gross¹; Etienne Dreyer¹

¹ *Weizmann Institute of Science (IL)*

Corresponding Authors: etienne.dreyer@cern.ch, eilam.work@gmail.com, nilotpal.kakati@cern.ch

Accurately reconstructing particles from detector data is a critical challenge in experimental particle physics, where the spatial resolution of calorimeters plays a key role. This study explores the integration of super-resolution techniques into the Large Hadron Collider (LHC)-like reconstruction pipeline to enhance the granularity of calorimeter data. By applying super-resolution, we demonstrate how significant improvements in reconstruction accuracy can be achieved without physical changes to the detectors. This approach could significantly impact the reconstruction pipeline of LHC-like experiments and could be a major consideration in future detector design.

Track:

Reconstruction

Detector Simulation / 73

CaloDREAM – Detector Response Emulation via Attentive flow Matching

Authors: Ayodele Ore^{None}; Luigi Favaro^{None}; Sofia Palacios Schweitzer¹; Tilman Plehn^{None}

¹ *ITP, University Heidelberg*

Corresponding Authors: luigi.favaro2@gmail.com, plehn@uni-heidelberg.de, ore@thphys.uni-heidelberg.de, palacios@thphys.uni-heidelberg.de

Detector simulations are an exciting application of modern generative networks. Their sparse high-dimensional data combined with the required precision poses a serious challenge. We show how combining Conditional Flow Matching with transformer elements allows us to simulate the detector phase space reliably. Namely, we use an autoregressive transformer to simulate the energy of each layer, and a vision transformer for the high-dimensional voxel distributions. We show how dimension reduction via latent diffusion allows us to train more efficiently and how diffusion networks can be evaluated faster with bespoke solvers. We showcase our framework, CaloDREAM, on datasets 2 and 3 of the CaloChallenge.

Track:

Detector simulation & event generation

Theorie / 74

Taming perturbation theory in QCD with Normalizing Flows

Author: Rikab Gambhir¹

Co-author: Radha Mastandrea²

¹ *MIT*

² *University of California, Berkeley*

Corresponding Authors: rikab@mit.edu, rmastand@berkeley.edu

When predicting the distribution of an observable, $p(x)$, in QCD, fixed-order (FO) perturbation theory can suffer from many undesirable artifacts, including large logarithms spoiling the expansion, unphysical divergences or negative bins, non-smooth kinks, and non-normalizability on physical x 's. However, one expects the “true” $p(x)$, as accessed by experiment, to be finite, positive, smooth, and normalized. We show how these conditions on $p(x)$ can be enforced exactly by parameterizing it using a Normalizing Flow (NF), which is matched onto FO calculations in regions of x where perturbation theory is expected to converge, which results in a “more physical” $p(x)$ that still agrees with perturbation theory. This performs an effective resummation of higher-order terms in taming divergences, constrained at the lowest orders to the perturbative expansion by the choice of loss, of which the usual leading logarithmic resummation is one possibility. In principle, additional physical structure including scheme independence, RG evolution (including DGLAP), factorization, or other constraints can be incorporated into the NF.

Track:

Theory

Event generation / 75

Exploring phase space with Flow Matching

Author: Timo Janssen^{None}

Corresponding Author: timo.janssen@theorie.physik.uni-goettingen.de

Generative models can speed up parton-level Monte Carlo event generation. Normalizing Flows are especially interesting due to their exact likelihood evaluation. Compared to discrete, layer-based flows, continuous Normalizing Flows (CNFs) have been shown to offer higher expressivity. New simulation-free training methods reduce their training costs significantly. We show that CNFs trained by Flow Matching can improve the sampling of parton-level QCD scattering events compared to traditional methods such as Vegas, both in terms of Monte Carlo variance and unweighting efficiency. We evaluate their performance for relevant LHC processes and compare it to discrete flows and Vegas.

Track:

Theorie / 76

Efficient SMEFT fits with neural importance sampling

Authors: Theo Heimel¹; Tilman Plehn^{None}; Nikita Schmal^{None}

¹ *Heidelberg University*

Corresponding Authors: schmal@thphys.uni-heidelberg.de, heimel@thphys.uni-heidelberg.de, plehn@uni-heidelberg.de

Global SMEFT analyses have become a key interpretation framework for LHC physics, quantifying how well a large set of kinematic measurements agrees with the Standard Model. We show how normalizing flows can be used to accelerate sampling from the SMEFT likelihood. The networks are trained without a pre-generated dataset by combining neural importance sampling with Markov chain methods. Furthermore, we use GPUs for fast evaluation of the likelihood, and compute profile likelihoods efficiently using differentiability.

Track:

Theory

Tagging / 78

WOTAN: Weakly-supervised Optimal Transport Attention-based Noise Mitigation

Author: Nathan Suri Jr.¹

Co-authors: Ben Nachman²; Vinicius Massami Mikuni²

¹ *Yale University (US)*

² *Lawrence Berkeley National Lab. (US)*

Corresponding Authors: benjamin.philip.nachman@cern.ch, vinicius.massami.mikuni@cern.ch, nathan.suri.jr@cern.ch

We improve upon the existing literature on pileup mitigation techniques studied at Large Hadron Collider (LHC) experiments for disentangling proton-proton collisions. Pileup presents a salient problem that, if not checked, hinders the search for new physics and Standard Model precision measurements such as jet energy, jet substructure, missing momentum, and lepton isolation. The primary technique that serves as the foundation for this work is known as Training Optimal Transport using Attention Learning (TOTAL). The TOTAL methodology compares matched samples with and without pileup interactions present to robustly learn an accurate description of pileup as a transport function without any need for assumptions of pileup nature derived from simulations. In this

work, we develop an improved version of TOTAL known as Weakly-supervised Optimal Transport Attention-based Noise Mitigation (WOTAN) by reducing the degree of TOTAL's self-supervision. The reduction in self-supervision allows us to demonstrate the power of optimal transport-based pileup mitigation in being able to use data for particle classification instead of solely simulations. Despite its reduced supervision, our work still outperforms existing conventional pileup mitigation approaches by improving the resolution of key observables relevant for both precision measurements and BSM searches in events with pileup interaction counts up to 200. WOTAN is the first fully data-driven machine learning pileup mitigation strategy capable of operating at LHC experiments.

Track:

Tagging (Classification)

Unfolding & Inference / 79

Advanced techniques for SBI in collider physics

Author: Giovanni De Crescenzo¹

¹ *University of Heidelberg*

Corresponding Author: giovanni.de_crescenzo@stud.uni-heidelberg.de

We present an application of Simulation-Based Inference (SBI) in collider physics, aiming to constrain anomalous interactions beyond the Standard Model (SM). This is achieved by leveraging Neural Networks to learn otherwise intractable likelihood ratios. We explore methods to incorporate the underlying physics structure into the likelihood estimation process. Specifically, we compare two approaches: morphing-aware likelihood estimation and derivative learning. Furthermore, we illustrate how uncertainty-aware networks can be employed to compare the performance of these methods. Additionally, we demonstrate two new techniques for enhancing the accuracy and reliability of the network training. First, we introduce of a new way to treat the outliers in the target reconstruction-level distributions by repeated smearing through a modified reweighting procedure (dubbed fractional smearing). Second, we utilise Lorentz-equivariant network architectures to exploit the symmetry structure inherent in the underlying particle physics amplitudes.

Track:

Unfolding

Reconstruction / 80

Accelerating Graph-based Tracking Tasks with Symbolic Regression

Author: Nathalie Soybelman¹

Co-authors: Carlo Schiavi²; Eilam Gross¹; Francesco Armando Di Bello²

¹ *Weizmann Institute of Science (IL)*

² *INFN e Universita Genova (IT)*

Corresponding Authors: francesco.armando.di.bello@cern.ch, nathalie.soybelman@cern.ch, carlo.schiavi@cern.ch, eilam.work@gmail.com

Reconstructing particle tracks from detector hits is computationally intensive due to the large combinatorics involved. Recent work has shown that ML techniques can enhance conventional tracking methods, but complex models are often difficult to implement on heterogeneous trigger systems, such as FPGAs. While deploying neural networks on FPGAs is possible, resource limitations pose challenges. As an alternative, we propose using symbolic regression (SR) to replace graph-based neural networks. This approach maintains the graph structure and enables message passing, making it more suitable for heterogeneous hardware. SR is easier to implement on FPGAs and offers faster execution on CPUs compared to traditional methods. Though demonstrated for tracking, this method provides a proof-of-concept applicable to various use cases.

Track:

Reconstruction

Detector Simulation / 81**Higher Resolution and Angular Conditioning for Normalizing-Flow-based Generation of Calorimeter Showers****Authors:** Anatolii Korol¹; Frank-Dieter Gaede¹; Gregor Kasieczka²; Peter McKeown³; Thorsten Lars Henrik Buss⁴¹ *Deutsches Elektronen-Synchrotron (DE)*² *Hamburg University (DE)*³ *CERN*⁴ *Universität Hamburg (DE)***Corresponding Authors:** peter.mckeown@cern.ch, anatolii.korol@desy.de, frank-dieter.gaede@cern.ch, gregor.kasieczka@cern.ch, thorsten.buss@desy.de

Monte Carlo (MC) simulations are crucial for collider experiments, enabling the comparison of experimental data with theoretical predictions. However, these simulations are computationally demanding, and future developments, like increased event rates, are expected to surpass available computational resources. Generative modeling can substantially cut computing costs by augmenting MC simulations, thereby addressing this issue.

To this end, we presented ConvL2LFlows, a convolutional-flow-based generative model, at last year's ML4Jets. This year, we present several improvements to this model, making it usable in realistic simulations. These improvements are: i) adding angular conditioning to generate showers with arbitrary incident angles ii) using nine times more bins than calorimeter readout-cells to be able to use the model for arbitrary incident points, and iii) integrating L2LFlows into the full simulation pipeline using DDFastShowerML.

We will systematically compare ConvL2LFlows with nine times higher resolution, ConvL2LFlows with cell-level granularity, and a point-cloud-based generative model named CaloClouds II. While fixed grid models like ConvL2LFlows represent showers as three-dimensional arrays, point cloud models represent them as unordered sets of points. Our comparison will highlight the advantages and disadvantages in a realistic setting.

Track:

Detector simulation & event generation

Tagging / 82

Hadronic Top Quark Polarimetry with ParticleNet

Authors: Alberto Navarro¹; Andrew Larkoski²; K.C. Kong^{None}; Zhongtian Dong³; dorival Gonalves¹

¹ *Oklahoma State University*

² *UCLA*

³ *University of Kansas*

Corresponding Authors: dorival@okstate.edu, alberto.navarro_serratos@okstate.edu, cdong@ku.edu, larkoski@g.ucla.edu, kckong@ku.edu

Observables sensitive to top quark polarization are important for characterizing and discovering new physics. The most powerful spin analyzer in the top decay is the down-type fermion from the W, which in the case of leptonic decay allows for very clean measurements. However, in many applications, it is useful to measure the polarization of hadronically decaying top quarks via an optimal hadronic spin analyzer. In this talk, we introduce and use subject flavor tagging to significantly improve spin analyzing power in hadronic decays beyond exclusive kinematic information employed in previous studies. We provide parametric estimates of the improvement from flavor tagging with any set of measured observables and demonstrate this in practice on simulated data using a Graph Neural Network (GNN).

Track:

Tagging (Classification)

Detector Simulation / 83

CaloClouds III: Ultra-Fast Geometry-Independent Highly-Granular Calorimeter Simulation

Authors: Anatolii Korol^{None}; Erik Buhmann¹; Frank-Dieter Gaede²; Gregor Kasieczka¹; Henry Day-Hall³; Katja Kruger²; Peter McKeown⁴; William Korcari¹

¹ *Hamburg University (DE)*

² *Deutsches Elektronen-Synchrotron (DE)*

³ *DESY*

⁴ *CERN*

Corresponding Authors: henry.day-hall@desy.de, peter.mckeown@cern.ch, william.korcari@cern.ch, erik.buhmann@desy.de, anatolii.korol@desy.de, gregor.kasieczka@cern.ch, frank-dieter.gaede@cern.ch, katja.kruger@cern.ch

Ever-increasing collision rates place significant computational stress on the simulation of future experiments in high energy physics. Generative machine learning (ML) models have been found to speed up and augment the most computationally intensive part of the traditional simulation chain: the calorimeter simulation. Many previous studies relied on fixed grid-like data representation of electromagnetic showers, which leads to artifacts when applied to highly granular calorimeters due to the aperiodic tiling of cells in realistic detector geometry. With this contribution, we present CaloClouds III, an updated version of the novel point cloud diffusion model, CaloClouds II. This new version features a simplified architecture that further accelerates inference time, along with added angular conditioning, allowing integration into the simulation pipeline. The model was tested in a realistic DD4hep based simulation model of the ILD detector concept for a future Higgs factory. This is done with the DDFastShowerML library which has been developed to allow for easy integration of generative fast simulation models into any DD4hep based detector model. With this it is possible to benchmark the performance of a generative ML model using fully reconstructed physics events by comparing them against the same events simulated with Geant4, thereby ultimately judging the fitness of the model for application in an experiment's Monte Carlo.

Track:

Detector simulation & event generation

Event generation / 84

Fast simulation of backgrounds at LHCb - a generalised tool

Author: Alex Marshall¹

¹ *University of Bristol (GB)*

Corresponding Author: am13743@bristol.ac.uk

Background estimation is already a bottleneck in several analyses at LHCb, and with the upcoming larger datasets, the demand for efficient background simulation will continue to grow. While there are existing tools that can provide quick, rough estimates of background reconstructed distributions (e.g. RapidSim), these cannot account for the effects of common selection criteria. The tool presented here addresses this limitation by utilising Variational Autoencoders (VAEs) to model the reconstruction and vertexing algorithms of LHCb. For any given decay channel this tool will generate tuples containing the same high-level variables as produced by the full LHCb simulation software. Analysts can use these generated tuples as if they were indeed produced by the full simulation, for example, one can apply any bespoke selection criteria—such as event filtering and MVA classifiers—and directly assess how that selection affects any given background. The tool can be used to quickly generate reliable fit component templates that can be used in analyses without requiring the computationally intensive LHCb simulation software.

Track:

Detector simulation & event generation

Event generation / 87

(R) Application of generative models for full-detector, whole-event simulated event generation and jet background subtraction

Co-authors: Brett Viren ¹; Dmitrii Torbunov ; Haiwang Yu ; Jin Huang ²; Meifeng Lin ; Shuhang Li ; Timothy Thomas Rinn ¹; Yeonju Go ³; Yi Huang ; Yihui Ren ³

¹ *Brookhaven National Laboratory*

² *Brookhaven National Lab*

³ *Brookhaven National Laboratory (US)*

Corresponding Authors: dtorbunov@bnl.gov, bv@bnl.gov, yren@bnl.gov, hyu@bnl.gov, sl4859@columbia.edu, yeonju.go@cern.ch, jhuang@bnl.gov, yhuang2@bnl.gov, mlin@bnl.gov, timothytrinn@gmail.com

AI generative models, such as generative adversarial networks (GANs), have been widely used and studied as efficient alternatives to traditional scientific simulations like Geant4. Diffusion models, which have demonstrated great capability in generating high-quality text-to-image translations in industry, have yet to be applied in the high-energy heavy-ion physics.

In this talk, we present the effectiveness of denoising diffusion probabilistic models (DDPMs) as AI-based generative surrogate models for whole-event, full-detector simulations in high-energy heavy-ion experiments [1]. We use HIJING minimum-bias data simulated by Geant4 with the sPHENIX geometry to train the model. We compare its performance with that of a popular alternative—

GANs. The results show that DDPMs significantly outperform GANs, providing much faster generation times compared to Geant4 simulations, with a speedup on the order of 100. This suggests the potential for DDPMs in accelerating complex event simulations in high energy collider experiments.

Additionally, unpaired image-to-image translation models can be applied for jet background subtraction techniques. We present that UVCGAN [2], one of the CycleGAN models, demonstrates excellent performance in separating jets from combinatorial background in heavy-ion collisions at both Relativistic Heavy Ion Collider and the Large Hadron Collider.

[1] Y. Go and D. Torbunov et al, Effectiveness of denoising diffusion probabilistic models for fast and high-fidelity whole-event simulation in high-energy heavy-ion experiments, <https://link.aps.org/doi/10.1103/PhysRevD.107.054001>, <https://arxiv.org/abs/2406.01602>

[2] D. Torbunov et al, UVCGAN v2: An Improved Cycle-Consistent GAN for Unpaired Image-to-Image Translation, <https://arxiv.org/abs/2303.16280>

Track:

Detector simulation & event generation

Unfolding & Inference / 88

Gamma-ray spectrometry of fission fragments : ML analysis of multi-dimensional spectra

Author: Mattéo Ballu^{None}

Corresponding Author: matteo.ballu@cea.fr

The analysis of gamma radiation emitted by fission fragments has become an essential tool for studying the nuclear fission process. It allows probing the intrinsic properties of the fragments or exploring effects that are little studied experimentally, such as the sharing of excitation energy between fragments during nuclear fission.

However, the analysis of experimental fission gamma-ray data using traditional techniques is time-consuming and complex. The main task is to find and extract peak intensities on 2D or 3D distributions (gamma-ray energies measured in coincidence), which are filled with thousands of peaks of variable amplitude, often overlapping with significant background noise. Classical methods rely on large models that can be difficult to fit.

To overcome this, we implemented a Convolutional Neural Network (UNET-like architecture) and trained it using synthetic data that closely imitate experimental data. To account for uncertainties in the input histograms and provide uncertainty estimates for the predicted intensities, we use an approach based on resampling and ensemble methods.

Preliminary results of applying the neural network to synthetic data indicate promising accuracy in identifying peak intensities, but further investigation is required to determine if this approach outperforms classical fit methods.

The final goal is to apply the trained model to real data obtained with the FIPPS instrument (a high-resolution HPGe spectrometer) at the nuclear facility of the Laue-Langevin Institute (ILL) to provide experimental verification of fission-delayed gamma-ray modelling.

Track:

Tagging / 89

Heavy-Flavour Jet Tagging at LHCb Using Graph Neural Networks

Authors: Gabriella Monet Pesticci¹; Nathan Grieser²

¹ CERN

² University of Cincinnati (US)

Corresponding Authors: nathan.allen.grieser@cern.ch, gabriella.monet.pesticci@cern.ch

Efficient jet flavour-tagging is crucial for event reconstruction and particle analyses in high energy physics (HEP). Graph Neural Networks (GNNs) excel in capturing complex relationships within graph-structured data, and we aim to enhance the classification of b-jets using this method of deep learning. Presented in this work is the first application of a novel GNN b-jet tagger using the LHCb detector, and plans for further expansion into different jet architectures. The fully-connected graphs are built using the different daughter particles associated to the jet as nodes, and global information from jet kinematics is used to improve performance. Beyond normal tracking information, the GNN presented makes use of LHCb's excellent abilities to have particle identification (PID) of the daughter particles, further enhancing the performance of the classification.

Track:

Tagging (Classification)

Uncertainties & Interpretability / 91

The good, the bad, and the Bayesian?

Authors: Luigi Favaro^{None}; Manuel Haußmann¹; Nina Elmer^{None}; Ramon Winterhalder²; Tilman Plehn^{None}

¹ Universität Heidelberg

² Università degli Studi di Milano

Corresponding Authors: rpw@mailbox.org, n.elmer@thphys.uni-heidelberg.de

Estimating uncertainties is a fundamental aspect in every physics problem, no measurements or calculations comes without uncertainties. Hence it is crucial to consider the effect of training a neural network to problems in physics. I will present our work on amplitude regression, using loop amplitudes from LHC processes, as an example to examine the impact of different uncertainties on the outcome of the network. We test the behavior of different neural networks with uncertainty estimation, including Bayesian neural networks and repulsive ensembles.

Track:

Uncertainties

Tagging / 92

The Fundamental Limit of Jet Tagging

Authors: Alexander Mück¹; Ben Nachman²; Humberto Reyes-González³; Joep Geuskens^{None}; Michael Kramer⁴; Nishank Nilesh Gite²; Vinicius Massami Mikuni²

¹ RWTH Aachen University

² *Lawrence Berkeley National Lab. (US)*

³ *RWTH Aachen*

⁴ *Rheinisch Westfaelische Tech. Hoch. (DE)*

Corresponding Authors: mueck@physik.rwth-aachen.de, nishank.nilesh.gite@cern.ch, vinicius.massami.mikuni@cern.ch, harg.zepelin@gmail.com, joep.geuskens@rwth-aachen.de, benjamin.philip.nachman@cern.ch, mkraemer@physik.rwth-aachen.de

Identifying the origin of high-energy hadronic jets ('jet tagging') has been a critical benchmark problem for machine learning in particle physics. Jets are ubiquitous at colliders and are complex objects that serve as prototypical examples of collections of particles to be categorized. Over the last decade, machine learning-based classifiers have replaced classical observables as the state of the art in jet tagging. Increasingly complex machine learning models are leading to increasingly more effective tagger performance. Our goal is to address the question of convergence - are we getting close to the fundamental limit on jet tagging or is there still potential for computational, statistical, and physical insight for further improvements? We address this question using state-of-the-art generative models to create a realistic, synthetic data with a known optimum. Various state-of-the-art taggers are deployed on this dataset, showing that there is a significant gap between their performance and the optimum. Our dataset and software are made public to provide a benchmark task for future developments in jet tagging and other areas of particle physics.

Track:

Tagging (Classification)

Reconstruction / 94

Graph Neural Network-Based Track Finding in the LHCb Vertex Detector

Authors: Anthony Correia¹; Fotis Giasemis¹; Nabil Garroum¹; Vava Gligorov¹

¹ *Centre National de la Recherche Scientifique (FR)*

Corresponding Authors: anthony.correia@cern.ch, vladimir.gligorov@cern.ch, fotis.giasemis@cern.ch, nabil.garroum@cern.ch

The next decade will see an order of magnitude increase in data collected by high-energy physics experiments, driven by the High-Luminosity LHC (HL-LHC). The reconstruction of charged particle trajectories (tracks) has always been a critical part of offline data processing pipelines. The complexity of HL-LHC data will however increasingly mandate track finding in all stages of an experiment's real-time processing. This paper presents a GNN-based track-finding pipeline tailored for the Run 3 LHCb experiment's vertex detector and benchmarks its physics performance and computational cost against existing classical algorithms on GPU architectures. A novelty of our work compared to existing GNN tracking pipelines is batched execution, in which the GPU evaluates the pipeline on hundreds of events in parallel. We evaluate the impact of neural-network quantisation on physics and computational performance, and comment on the outlook for GNN tracking algorithms for other parts of the LHCb track-finding pipeline.

Track:

Reconstruction

Unfolding & Inference / 96**Parameter Estimation with Neural Simulation-Based Inference in ATLAS****Co-authors:** Arnaud Jean Maury ¹; David Rousseau ²¹ *Université Paris-Saclay (FR)*² *IJCLab-Orsay***Corresponding Authors:** rousseau@lal.in2p3.fr, arnaud.maury@universite-paris-saclay.fr

Neural Simulation-Based Inference (NSBI) is a powerful class of machine learning (ML)-based methods for statistical inference that naturally handle high dimensional parameter estimation without the need to bin data into low-dimensional summary histograms. Such methods are promising for a range of measurements at the Large Hadron Collider, where no single observable may be optimal to scan over the entire theoretical phase space under consideration, or where binning data into histograms could result in a loss of sensitivity. This work develops an NSBI framework that, for the first time, allows NSBI to be applied to a full-scale LHC analysis, by successfully incorporating a large number of systematic uncertainties, quantifying the uncertainty coming from finite training statistics, developing a method to construct confidence intervals, and demonstrating a series of intermediate diagnostic checks that can be performed to validate the robustness of the method. As an example, the power and feasibility of the method are demonstrated for an off-shell Higgs boson couplings measurement in the four lepton decay channel, using ATLAS experiment simulated samples. The proposed method is a generalisation of the standard statistical framework at the LHC, and can benefit a large number of physics analyses. This work serves as a blueprint for measurements at the LHC using NSBI.

Track:

Unfolding

Anomaly detection / 98**Efficient Resonant Anomaly Detection****Authors:** David Shih^{None}; Ranit Das¹¹ *Rutgers University***Corresponding Authors:** dshih@physics.rutgers.edu, rd804@scarletmail.rutgers.edu

A key step in any resonant anomaly detection search is accurate estimation of the background distribution in each signal region. Data-driven methods like CATHODE accomplish this by training separate density estimators on the complement of each signal region, and interpolating them into their corresponding signal regions. Having to re-train the density estimator on essentially the entire dataset for each signal region is a major computational cost in a typical sliding window search with many signal regions. We present a new method which significantly reduces this computational cost, while retaining a similar high quality of background density estimation and sensitivity to anomalous signals.

Track:

Anomaly detection

Detector Simulation / 99**Refining CMS Fast Simulation with ML-based regression****Co-authors:** CMS Collaboration ; Samuel Louis Bein ¹¹ *Universite Catholique de Louvain (UCL) (BE)***Corresponding Author:** samuel.bein@cern.ch

The CMS Fast Simulation chain (FastSim) is roughly 10 times faster than the application based on the GEANT4 detector simulation and full reconstruction referred to as FullSim. This advantage however comes at the price of decreased accuracy in some of the final analysis observables. A machine learning-based technique to refine those observables has been developed and its status is presented here. We employ a regression neural network trained within the framework of Fast Perfekt, using a combination of multiple loss functions to provide post-hoc corrections to samples produced by the FastSim chain. This technique results in a higher accuracy FastSim and thus allows for wider usage of FastSim.

Track:

Detector simulation & event generation

Mixed contributions / 100**Introducing Aspen Open Jets: a real-world ML-ready dataset for jet physics****Authors:** Alexander Mück¹; Anna Maria Cecilia Hallin²; Darius Faroughy³; David Shih^{None}; Gregor Kasieczka⁴; Humberto Reyes-González⁵; Ian Pang^{None}; Joschka Birk⁴; Luca Anzalone^{None}; Michael Kramer⁶; Oz Amram⁷¹ *RWTH Aachen University*² *University of Hamburg*³ *University of Zurich*⁴ *Hamburg University (DE)*⁵ *RWTH Aachen*⁶ *Rheinisch Westfaelische Tech. Hoch. (DE)*⁷ *Fermi National Accelerator Lab. (US)***Corresponding Authors:** oz.amram@cern.ch, gregor.kasieczka@cern.ch, ian.pang@physics.rutgers.edu, dshih@physics.rutgers.edu, mkraemer@physik.rwth-aachen.de, anna.hallin@uni-hamburg.de, joschka.birk@cern.ch, luca.anzalone2@unibo.it, faroughy@physik.uzh.ch, mueck@physik.rwth-aachen.de, harg.zepelin@gmail.com

We present Aspen Open Jets, a dataset consisting of 170M unlabelled jets derived from the CMS Open Data 2016. We show how using this dataset in the context of pre-training a foundation model can reduce the need for expensive simulated datasets. The dataset includes event information, jet kinematics, jet tagging information, particle kinematics, displacement, charge, PID and PUPPI weights, and will be available for further use by the community.

Track:**Uncertainties & Interpretability / 101**

Uncertainty Quantification and Anomaly Detection with Evidential Deep Learning

Author: Mark Neubauer¹

¹ *Univ. Illinois at Urbana Champaign (US)*

Corresponding Author: msn@illinois.edu

Evidential Deep Learning (EDL) is an uncertainty-aware deep learning approach designed to provide confidence (or epistemic uncertainty) about test data. It treats learning as an evidence acquisition process where more evidence is interpreted as increased predictive confidence. This talk will provide a brief overview of EDL for uncertainty quantification (UQ) and its application to jet tagging in HEP. I will also discuss connections between UQ and anomaly detection (AD) to describe some on-going work on improved AD using EDL methods.

Track:

Event generation / 102

Generating particle-clouds with discrete features using Markov jump processes

Author: Darius Faroughy¹

¹ *Rutgers University*

Corresponding Author: darius.faroughy@rutgers.edu

In many real-world scenarios, data is hybrid —i.e. described by both continuous and discrete features. At high-energy accelerators like the LHC, jet constituents exhibit discrete properties such as electric charge or particle-id. In this talk, we introduce a novel generative model for discrete features based on continuous-time Markov jump processes. By combining our approach with well-known models for continuous features, such as diffusion or flow-matching, we can effectively model hybrid data within a unified framework. We apply our method to generate particle-clouds that incorporate kinematic data, particle identities, and charge information. We demonstrate the effectiveness of our approach on the JetClass dataset.

Track:

Detector simulation & event generation

Tagging / 103

Transforming Flavour Tagging on ATLAS

Authors: Dan Guest¹; Greta Brianti²; Nicole Michelle Hartman³; Samuel Van Stroud⁴; Valentina Vecchio⁵

¹ *Humboldt University of Berlin (DE)*

² *CERN*

³ *TUM (DE)*

⁴ *UCL*

⁵ *University of Manchester*

Corresponding Authors: greta.brianti@cern.ch, samuel.van.stroud@cern.ch, nicole.hartman@cern.ch, valentina.vecchio@cern.ch, daniel.hay.guest@cern.ch

Flavour-tagging is a critical component of the ATLAS experiment's physics programme. Existing flavour tagging algorithms rely on several 'low-level' taggers, which are a combination of physically informed algorithms and machine learning models. A novel approach presented here instead uses a single machine learning model based on reconstructed tracks, avoiding the need for low-level taggers based on secondary vertexing algorithms. This new approach reduces complexity and improves tagging performance. This model employs a transformer architecture to process information from a variable number of tracks and other objects in the jet in order to simultaneously predict the jet's flavour, the partitioning of tracks into vertices, and the physical origin of each track. The inclusion of auxiliary tasks aids the model's interpretability. The new approach significantly improves jet flavour identification performance compared to existing methods in both Monte-Carlo simulation and collision data. Notably, the versatility of the approach is demonstrated by its successful application in boosted Higgs tagging using large-R jets.

Track:

Tagging (Classification)

Anomaly detection / 104

Model-agnostic search for dijet resonances with anomalous jet substructure in proton-proton collisions at $\sqrt{s} = 13$ TeV with the CMS detector

Authors: Aritra Bal¹; Hongbo Liao²

¹ *KIT - Karlsruhe Institute of Technology (DE)*

² *Chinese Academy of Sciences (CN)*

Corresponding Authors: aritra.bal@cern.ch, hongbo.liao@cern.ch

We introduce a model-agnostic search for new physics in the dijet final state. Other than the requirement of a narrow dijet resonance with a mass in the range of 1.8-6 TeV, minimal additional assumptions are placed on the signal hypothesis. Search regions are obtained by utilizing multivariate machine learning methods to select jets with anomalous substructure. A collection of complementary anomaly detection methods – based on unsupervised, weakly-supervised and semi-supervised algorithms – are used in order to maximize the sensitivity to unknown new physics signatures. These algorithms are applied to data corresponding to an integrated luminosity of 137 inverse femtobarns, recorded in the years 2016 to 2018 by the CMS experiment at the LHC, at a centre-of-mass energy of 13 TeV. Exclusion limits are derived on the production cross section of benchmark signal models varying in resonance mass, jet mass and jet substructure. Many of these signatures have not previously been searched for at the LHC, making the limits reported on the corresponding benchmark models the first ever and the most stringent to date.

Track:

Anomaly detection

Uncertainties & Interpretability / 107

ParticleNet: Calibration of the jet energy scale using pT regression with partial Run3 data collected by the CMS experiment

Authors: Hongbo Liao¹; Matteo Malucchi²

¹ *Chinese Academy of Sciences (CN)*

² *ETH Zurich (CH)*

Corresponding Authors: matteo.malucchi@cern.ch, hongbo.liao@cern.ch

We are presenting the first calibration of the jet pT regression (CMS-DP-2024-064), achieving an expected improvement in jet resolution of up to 17%, and the latest performance results for flavor identification and jet energy resolution estimation using ParticleNet. The pT regression, which focuses on correcting the reconstructed jet pT to the truth-level jet pT, is divided into two components: the visible part due to detector and jet misreconstructions and the invisible part including the regression of the neutrinos which are not reconstructed by the detector. A key focus is represented by the jet energy scale calibration for the regressed pT, based on data from proton-proton collisions at $\sqrt{s} = 13.6$ TeV in 2022 and 2023. The results are shown for jets clustered from particle flow candidates using the anti-kT algorithm with radius parameter 0.4, and applying the Pileup Per Particle Identification algorithm for pileup mitigation. Our findings demonstrate that the standard jet correction chain can be successfully applied to the regressed pT.

Track:

Uncertainties

Tagging / 108

UParT: A unified approach for jet-based object identification in CMS in Run 3

Co-authors: Hongbo Liao¹; Uttiya Sarkar²

¹ *Chinese Academy of Sciences (CN)*

² *Rheinisch Westfaelische Tech. Hoch. (DE)*

Corresponding Authors: uttiya.sarkar@cern.ch, hongbo.liao@cern.ch

The steady progress in machine learning leads to substantial performance improvements in various areas of high-energy physics, especially for object identification. Jet flavor identification (tagging) is a prominent benchmark that profits from elaborate architectures, leveraging information from low-level input variables and their correlations. Throughout the data-taking eras of the Large Hadron Collider (LHC) (Run 1 - Run 3), various deep-learning-based algorithms were established and led to a significantly improved tagging performance of heavy flavor jets, originating from the hadronization of b and c quarks. Individual developments led to the extension of heavy-flavor jet tagging to hadronic τ jet identification and simultaneous jet energy regression. At the same time, using a so-called adversarial training strategy, the robustness of algorithms was increased, reducing the dependence on possible mismodeling in simulation compared to data. This note presents a new approach for object tagging based on jets, unifying different approaches and extending the paradigm of b and c jet identification to s and hadronic τ jet identification, simultaneous flavor aware jet energy and resolution regression, and incorporating an innovative adversarial training approach. We show that the new algorithm based on the ParticleTransformer architecture denoted UParT, presents an advantageous algorithm for Run 3.

Track:

Tagging (Classification)

Tagging / 109

Performance of the CNN-based tau identification algorithm with Domain Adaptation using Adversarial Machine Learning for Run 2

Authors: Hongbo Liao¹; Olha Lavoryk²

¹ *Chinese Academy of Sciences (CN)*

² *KIT - Karlsruhe Institute of Technology (DE)*

Corresponding Authors: olha.lavoryk@cern.ch, hongbo.liao@cern.ch

Precise tau identification is a crucial component for many studies targeting the Standard Model or searches for New Physics within the CMS physics program. The Deep Tau v2.5 algorithm is a convolutional neural network algorithm: an improved version of its predecessor, Deep Tau v2.1, deployed for the LHC Run 3. This updated version integrates several enhancements to improve classification performance, including domain adaptation techniques, expanded datasets, and improved feature standardization, hyperparameter optimization, and data balancing. These improvements make the model more robust against potential discrepancies between data and simulation that would lead to a bias in the training. The enhancements result in a notable improvement in accuracy, with Deep Tau 2.5 achieving approximately a 30% reduction in background contributions with respect to its predecessor

Track:

Tagging (Classification)

Reconstruction / 111

Deep learning on jet modification in the presence of the QGP background

Corresponding Author: ranli@mail.sdu.edu.cn

Jet interactions with the color-deconfined QCD medium in relativistic heavy-ion collisions are conventionally assessed by measuring the modification of the distributions of jet observables with respect to their proton-proton baselines. Deep learning methods allow us to evaluate the modification of jets on a jet-by-jet basis, and therefore significantly improve the capability of using jets to probe the QGP properties. In this work, we first explore the fractional energy loss of each jet through the QGP using the Convolutional Neural Network (CNN) method. The initial jets are generated by Pythia, and their subsequent evolution through the QGP is simulated using a linear Boltzmann transport (LBT) model that incorporates both elastic and inelastic scatterings between jet partons and the QGP. By mixing jet partons with the QGP background generated by a thermal model, and then training the neural network with jets obtained using the constituent subtraction method, we show the neural network can provide a good prediction on the fractional energy loss of jets in the presence of the QGP background. We further apply the Dense Neural Network (DNN) method and the aforementioned CNN method to the background subtraction in constructing jets. Although the recoil partons from the LBT simulation, scattered out of the QGP background but belonging to jets, can inevitably lead to over-subtraction of the background, we obtain better accuracy of background subtraction by using the deep learning methods than by using the traditional constituent subtraction method and area-based method adopted in many experimental measurements.

Track:

Reconstruction / 112

Transformer networks for constituent-based b-jet calibration with the ATLAS detector

Corresponding Author: brendon.aurele.bullard@cern.ch

The precise measurement of kinematic features of jets is key to the physics program of the LHC. The determination of the energy and mass of jets containing bottom quarks b -jets is particularly difficult given their distinct radiation patterns and production of undetectable neutrinos via leptonic heavy flavor decays. This talk will describe a novel calibration technique for the b -jet kinematics using transformer-based neural networks trained on simulation samples. Separate simulation-based regression methods have been developed to estimate the transverse momentum of small-radius jets and the transverse momentum and mass of large-radius jets. In both cases, the medians of reconstructed jet properties are corrected to the true value across a range of jet features. A relative jet energy resolution improvement with respect to the nominal calibration between 18% and 31% is demonstrated for small-radius jets. Both the large-radius jet transverse momentum and mass resolution are shown to improve by 25–35%. These methods improve meaningfully upon simulation-based b -jet correction strategies previously used in ATLAS.

Track:**Reconstruction / 113**

Jet Finding as a Real-Time Object Detection Task

Corresponding Author: leon.bozianu@cern.ch

The High Luminosity upgrade to the LHC will deliver an unprecedented luminosity to the ATLAS experiment. Ahead of this increase in data the ATLAS trigger and data acquisition system will undergo a comprehensive upgrade. The key function of the trigger system is to maintain a high signal efficiency together with a high background rejection whilst adhering to the throughput constraints of the data acquisition system. Here we propose a calorimeter-only fast preselection step to speed up the trigger decision for hadronic signals containing jets.

In this work we present the design and implementation of an object detection Convolutional Neural Network (CNN) for jet finding in the ATLAS calorimeter. The model is employed in the task of jet detection to identify and localise jets within the full calorimeter acceptance and to subsequently estimate their transverse momenta.

The performance of the object detection architecture, which targets real-time applications, is evaluated on a set of simulated particle interactions in the ATLAS detector with up to 200 concurrent pile-up interactions.

Reconstruction / 114

(R) Generative Neural Networks for Reconstructing Parton-Level Jet Showers after Hadronization

Corresponding Author: umar.sohail.queshi@vanderbilt.edu

Recreating realistic parton-level event configurations from jets is a crucial task for various physics analyses. However, hadronization processes cannot be computed using perturbative QCD. Therefore,

it has been traditionally intractable to reconstruct parton-level events after hadronization.

We present a generative machine learning approach for reconstructing jet showers at the parton level from hadron-level jets. In particular, we utilize state-of-the-art generative models and vector representations of jets $\mathcal{J} = \{n, (p_i^\mu, \eta_i, \phi_i)_{i=1}^n\}$, where n is the particle multiplicity. Unlike traditional regression-based methods that focus on predicting individual particle properties, our method captures the entire parton-level event structure from jet data, offering a physically realistic reconstruction.

For this talk, we look at jets originating from photon-tagged events to maximize partonic structure in a single reconstructed jet, although our method works for any jet multiplicity and process. We evaluate the performance of our method using the energy mover's distance metric, in addition to studying the impact of different sources of background such as underlying events, detector effects, and pileup events.

Track:

Reconstruction / 115

Particle flow and flavor tagging with DNN for Higgs factories

Corresponding Author: suehara@icepp.s.u-tokyo.ac.jp

Deep learning can give a significant impact on physics performance of electron-positron Higgs factories such as ILC and FCCee. We are working on two topics on event reconstruction to apply deep learning; one is jet flavor tagging. We apply particle transformer to ILD full simulation to obtain jet flavor, including strange tagging. The other one is particle flow, which clusters calorimeter hits and assigns tracks to them to improve jet energy resolution. We modified the algorithm developed in context of CMS HGCALE based on GravNet and Object Condensation techniques and add a track-cluster assignment function into the network. The overview and performance of these algorithms will be presented.

We believe the sophisticated simulation developed for long time in ILD context is essential to try these novel technologies in event reconstruction. Comparison with other Higgs factory results as well as primal consideration on impact to physics performance will also be discussed.

Mixed contributions / 117

ML assisted Event Reconstruction in the CMS Phase-2 High Granularity Calorimeter Endcap

Co-authors: Hongbo Liao¹; Theo Cuisset²

¹ Chinese Academy of Sciences (CN)

² LLR / École Polytechnique (FR)

Corresponding Authors: theo.cuisset@polytechnique.edu, hongbo.liao@cern.ch

The high-luminosity era of the LHC will pose unprecedented challenges to the detectors. To meet these challenges, the CMS detector will undergo several upgrades, including the replacement of the current endcap calorimeters with a novel High-Granularity Calorimeter (HGCALE). To make optimal use of this innovative detector, novel algorithms have to be invented. A dedicated reconstruction framework, The Iterative Clustering (TICL), is being developed within the CMS Software (CMSSW).

This new framework is designed to fully exploit the high spatial resolution and precise timing information provided by HGCal. Several key ingredients of the object reconstruction chain already rely on Machine Learning techniques and their usage is expected to further develop in the future. In the presentation, the reconstruction existing strategies will be presented stressing the role played by ML techniques to exploit the information provided by the detector. The areas where ML techniques are expected to play a role in the future developments will be also discussed

Track:

Reconstruction / 118

Synergizing Physics: Deep Learning Techniques for Time-of-Flight Reconstruction and Jet Tagging in High Energy Physics

Author: Konrad Helms¹

¹ *Georg-August-Universität Göttingen*

Corresponding Author: konrad.helms@theorie.physik.uni-goettingen.de

This talk presents a synergy between quark/gluon jet tagging on LHC data, and charged hadron time-of-flight (TOF) regression on ILC data, in the form of one problem-solving mechanism that can address both tasks. They both involve processing data represented as unordered point clouds of varying sequence lengths, optimally handled using permutation-invariant architectures.

A transformer-based quark/gluon jet classifier is introduced and compared to a convolution-based model, serving as a benchmark. Both networks operate on sets of jet constituents and per-jet observables. The transformer-based architecture outperforms its compared-to benchmark, and is on par with modern state-of-the-art models.

The transformer-based TOF estimator outperforms the current best estimator used by the ILD community, and, due to its remarkable lack of bias in flight time predictions, is optimally suited for time-of-flight-based mass estimation.

Track:

Plenary talks / 119

Symbolic machine learning in physics

Corresponding Author: fcharton@gmail.com

Transformers excel at symbolic data manipulation, but most of their applications in physics deal with numerical calculations. I present a number of applications of symbolic AI in mathematics, and one in theoretical physics: learning scattering amplitudes.

Track:

Plenary talks / 120

Experimental highlights: Edge AI for particle physics

Corresponding Author: jennifer.ngadiuba@cern.ch

The ever-growing data volumes produced by HEP experiments, particularly at the CERN Large Hadron Collider (LHC) and upcoming facilities, demand innovative approaches to data processing and analysis. Traditional data acquisition and processing methods are no longer adequate for handling the scale, speed, and complexity of this data. In response, the field has seen a transformative shift toward edge AI for intelligent trigger and front-end systems, fundamentally changing how experiments manage data acquisition and processing in real-time. This talk will cover promising implementations of these new approaches in current and future HEP experiments and their impact in accelerating discovery and pushing the boundaries of scientific knowledge in high-energy physics and beyond.

Track:

Plenary talks / 121

Experimental highlights: The State of ML in LHC Science

Corresponding Author: daniel.thomas.murnane@cern.ch

Informed by the many fields in which machine learning (ML) has made impacts, the coming years promise to see exciting improvements in the discovery and measurement power of LHC experiments. But stepping back from the many exploratory studies ongoing, there are already dozens of concrete and rigorous public LHC results leveraging advanced ML. This review will examine common themes of those results across simulation, offline reconstruction, BSM searches and precision analysis, and simulation-based inference. We thus find experiments converging towards similar techniques for some applications and diverging on others, with longstanding challenges in production-ready ML addressed even as others arise.

Track:

Plenary talks / 122

MadNIS - A journey towards the first ML event generator

Corresponding Author: rpw@mailbox.org

High-precision simulations based on first principles are a cornerstone of LHC physics research. In view of the HL-LHC era, there is an ever-increasing demand for both accuracy and speed in simulations. In this talk, I will first explain the basic principles of LHC event generation and highlight current methodologies and their bottlenecks. Afterwards, I will delve into the MadNIS journey and illustrate how modern ML techniques, as well as advanced computing hardware, can alleviate these limitations. In particular, I will present recent advancements in neural importance sampling, the development of fast amplitude surrogates, and the latest progress on GPU-accelerated MadGraph.

Track:

Plenary talks / 123

Implicit and Explicit Simulation-Based Inference for Cosmology

Corresponding Author: francois.lanusse@cea.fr

Physical models in the form of simulations offer an avenue to model the data in all of its complexity, but until very recently using such models to estimate physical fields and parameters remained an open problem.

In this talk, I will discuss two possible points of view on simulators, depending on whether they are “black-box” or “open-box” models, and the different methodologies and strategies which may be applied in each case to use these physical models within a Bayesian inference context. As both of these approaches become tractable, an interesting question for our field is to discuss which point of view will be the most effective and robust in practice.

In the case of black-box simulations (which can only be sampled from), I will discuss in particular considerations of optimal data compression for simulation-based inference.

In the case of open-box simulations, which can be seen as differentiable probabilistic models, with an explicit joint log probability, I will discuss strategies and challenges for building large scale differentiable physical models of the Universe touching in particular on distributed differentiable N-body solvers and building accelerated hybrid physical/ml simulations leveraging neural ODE methodologies.

Track:

Plenary talks / 124

Theory Overview

Corresponding Author: sven.krippendorf@lmu.de

In this talk I will give a biased review on the work at the intersection of machine learning and theoretical physics. This includes how we can use transformers to obtain symbolic expressions without having information about the target expression. In turn, I present a benchmark human physicists have failed in solving, namely that of compact Calabi-Yau metrics and give a short status report on ML attempts. I then discuss how we can use efficient use of automatic differentiation to enable for the first time the large-scale exploration of string theory solutions. To round it off I shortly comment on how we can use automated theorem proving to formalize certain questions in theoretical physics and how theoretical physics can help us building more efficient neural networks.

Track:

Plenary talks / 125

Realtime reconstruction: Machine learning in reconstruction at LHC

Corresponding Author: simon.akar@cern.ch

The Large Hadron Collider (LHC) at CERN pushes the boundaries of particle physics, generating data at unprecedented rates and requiring advanced computational techniques to process information in real time. While experimental environments between LHC experiments can differ, common challenges can be identified in the area of real-time reconstruction including the use of specialized

trigger systems, machine learning techniques for fast data reduction, and GPU/FPGA-accelerated architectures that allow efficient processing within microsecond latencies. This talk will cover a subset of the most recent ML applications for realtime reconstruction at the LHC experiments.

Track:

Plenary talks / 126

Time Series Anomaly Detection: Overview and New Trends

Corresponding Author: boniol.paul@gmail.com

Anomaly detection is an important problem in data analytics with applications in many domains. In recent years, there has been an increasing interest in anomaly detection tasks applied to time series. In this talk, we take a holistic view of anomaly detection in time series, discussing the challenges and research opportunities in this field. In addition, we will focus on the challenges related to anomaly detection in heterogeneous time series datasets, as well as on the new research opportunities related to Model Selection and Ensembling.

Welcome & Opening Talks / 127

Welcome

Closing remarks / 128

Closing remarks

Corresponding Author: abutter@lpnhe.in2p3.fr