

Artificial Intelligence and the Uncertainty challenge in Fundamental Physics

lundi 27 novembre 2023 - vendredi 1 décembre 2023



Recueil des résumés

Contents

Using an adversary trained on a control sample to control systematic errors	1
An unfolding method based on conditional Invertible Neural Networks (cINN) using iterative training	1
Can contrastive learning de-bias my model?	1
Metrics for Uncertainty-Aware ML Methods	2
Advancing Explainable AI: Testing and Enhancing Techniques Across Multidisciplinary Use-Cases	2
Precision-Machine Learning for the Matrix Element Method	3
Uncertainty-aware diffusion models for LHC Event Generation	3
Publication and reuse of ML models in LHC analyses	4
Data driven background estimation in HEP using generative adversarial networks	4
Swyft: Direct marginal inference for large simulation models	4
Generative modeling in genomics and a perspective on uncertainty quantification	5
Efficient Sampling from Bayesian Network Posteriors for Optimal Uncertainties	5
Uncertainty Quantification in Neural Networks: Methods and Considerations	6
Optimal dataset-wide inference in the presence of systematic uncertainties	6
Exploring Data Challenges and Leveraging Codabench: A Practical Journey with unsupervised New Physics detection at 40 MHz	6
Neuroscience ML challenges using CodaBench: Decoding multi-limb trajectories from two-photon calcium imaging	7
A Data-centric AI based diagnosis method using unsupervised labeling and MLP	7
Data frugal machine learning approaches for unmixing problems in Physics	8
Potential and challenges of highly dimensional generative models	8
Robustness to Uncertainties in Machine Learning Applications for HEP	8

Machine-learning and equations-informed tools for generation and augmentation of turbulent data.	9
Fair Universe Challenge Paris 2023	9
Revisiting models and uncertainty with AI	10
Data Subsampling for Bayesian Neural Networks	10
Fair Universe HiggsML Uncertainty Challenge: Lessons Learned and plans	11
Overview of HiggsML Uncertainty Challenge	11
Uncertainty Quantification in High Energy Physics	11
Bayesian optimisation	11
Uncertainty Quantification in Industry	11
Welcome	12
Simulation-Based Inference: Where Classical Statistics Meets Machine Learning	12
Explainable AI for Interpretability of Deep Neural Networks : the High Energy Physics perspective	12
xAI in practice: current state of the art, limitations and perspectives	13
SBI	13
Simulation Based Inference: The Frequentist Perspective Abstract:	13
SBI Introduction	13
TBC	13
Unfolding in High Energy Physics	13
Uncertainty Quantification and Anomaly Detection with Evidential Deep Learning	14
Highlights from the 1st CNRS AISSAI Thematic Quarter on Causality	14
Machine Learning Assisted Sampling: Applications in Physics	14
Fair Universe hackathon : outcome and plans	15
Farewell	15
Uncertainty modeling in particle physics	15
An Introduction to Bayesian Optimization	15
Walkthrough of CodaBench and submissions	16
Scoring and Baseline Systematic Aware method	16
Feedback from participants	16

Hands-on hackathon and discussion	16
Wrap up and next steps	16
Uncertainty Quantification and Machine Learning in Industry: Current Practices and Challenges in Industrial Applications for Low-Carbon Electricity Production	17
An introduction to simulation-based inference	17
Introduction at Institut Pascal	17
Introduction	17

Controlling uncertainties in generative models / 1**Using an adversary trained on a control sample to control systematic errors****Auteur:** Gordon Watts¹¹ *University of Washington*

Machine Learning improved the sensitivity in searches for massive long-lived neutral particles decaying in the Calorimeter by over 30%. This was only after suppressing a large increase in the systematic errors caused by the method. The largest contribution to this improvement in sensitivity is the use of a Recurrent Neural Network that separates signal from standard QCD multijet background and Beam Induced Background. This classifier uses low-level data like realtive calorimeter cluster locations, tracks, and muon segments. We exploit the calorimeter cell energy deposit time as a powerful handle to reject beam induced background, which is poorly simulated by the ATLAS experiment's Monte Carlo simulation package. In addition, the beam induced background training dataset can only be drawn from data. Thus the RNN training set contains a mix of poorly simulated Monte Carlo data, and LHC collision data. A control dataset was used to train an adversary, along with signal and background samples for the RNN, simultaneously. The adversary is trained to tell the difference between collision data and simulated data, and its success is part of the main network's loss function. This dramatically reduced the systematic errors due to Monte Carlo mis-modeling. This presentatino will discuss the network design, how it was modified when the problem(s) were discovered, and its performance.

Unfolding (de-biasing) / 2**An unfolding method based on conditional Invertible Neural Networks (cINN) using iterative training****Auteurs:** Anja Butter¹; Bogdan MALAESCU²; Mathias Josef Backes³; Monica Dunford³¹ *LPNHE*² *LPNHE, Paris, FRANCE*³ *Kirchhoff Institut für Physik*

The unfolding of detector effects is crucial for the comparison of data to theory predictions. While traditional methods are limited to representing the data in a low number of dimensions, machine learning has enabled new unfolding techniques while retaining the full dimensionality. Generative networks like invertible neural networks (INN) enable a probabilistic unfolding, which maps individual data events to their corresponding unfolded probability distribution. The accuracy of such methods is however limited by how well the experimental data is modeled by the simulated training samples.

We introduce the iterative conditional INN (IcINN) for unfolding that adjusts for deviations between simulated training samples and data. The IcINN unfolding is first validated on toy data and then applied to pseudo-data for the $pp \rightarrow Z\gamma\gamma$ process. Additionally, we validate the probabilistic unfolding with a novel approach using the traditional transfer matrix-based methods.

The main results of this project have been published in a paper (arXiv:2212.08674: <https://arxiv.org/abs/2212.08674>). A second paper with a stronger focus on the probabilistic unfolding will be published prior to the conference.

Unfolding (de-biasing) / 3

Can contrastive learning de-bias my model?

Auteurs: Alex Wilkinson¹; Radi Radev²

¹ *UCL / Fermilab*

² *CERN*

Deep learning models have become ubiquitous in high-energy physics and have been successfully applied to a wide variety of tasks. Models for reconstruction are usually trained from scratch on a nominal set of simulation parameters, not taking into account variations of detector systematic uncertainties.

Following advances in contrastive learning, we present a method of pre-training a general model, that is de-biased from detector systematic uncertainties. During the pre-training phase, the model learns a representation that is invariant to simulation shifts and symmetries, by contrasting between different simulated views of the same event. Freezing the weights of the contrastive model, the extracted representation is general enough that it can be used for a variety of prediction tasks. We showcase the flexibility and efficacy of this method by training with sparse 3D neutrino liquid argon time projection chamber (LArTPC) data.

Opening session, Uncertainty Quantification / 4

Metrics for Uncertainty-Aware ML Methods

Auteurs: Benjamin Nachman¹; Sascha Diefenbacher¹; Wahid Bhimji¹

¹ *Lawrence Berkeley National Laboratory*

Machine learning methods have managed to provide significant improvements to data analysis in a multitude of scientific fields. However, as ML finds more and more applications in science, the challenge of quantifying machine learning uncertainties moves into the forefront. This is especially notable in High Energy Physics, where high-precision measurements require precise knowledge of uncertainties. Moreover, systematic uncertainties, such as detector effects and calibration factors, are common occurrences in HEP.

This leads to a requirement for methods and approaches that are not only accurate and precise in the presence of such, imperfectly understood systematic uncertainties, but can also provide accurate estimates of the uncertainty in their prediction. Several methods have been proposed for ML uncertainty quantification, however measuring and comparing the performance of these methods is highly non-trivial.

In this talk, we present several metrics for uncertainty quantification metrics, compare their distinct advantages, and benchmark them with example uncertainty quantification challenges.

Explainable AI / 5

Advancing Explainable AI: Testing and Enhancing Techniques Across Multidisciplinary Use-Cases

Auteurs: Andrea Ciardiello¹; Cecilia Voena¹; Corneliu Balan²; Cristiano Sebastiani³; Joseph Carmignani³; Kalina Dimitrova⁴; Maurizio Mattia⁵; Monica D'Onofrio³; Simone Melchionna⁶; Simone Scardapane⁷; Stefano Giagu⁸; Venelin Kozhuharov⁹

¹ *Sapienza Università di Roma, Roma, Italy*

² *Polytechnic Institute Bucharest Romania*

³ *University of Liverpool*

⁴ *Sofia University "St. Kliment Ohridski"*

⁵ *Istituto Superiore di Sanità - Roma, Italy*

⁶ *MedLea s.r.l.s.*

⁷ *Sapienza University of Rome, Roma, Italy*

⁸ *Sapienza Università di Roma, INFN Roma, Roma Italy*

⁹ *Sofia University "St. Kl. Ohridski"*

Developing and testing methodologies for enhancing the transparency, interpretability, and explainability of AI algorithms is a pressing challenge for the application of artificial intelligence methods in fundamental physics. The Multi-disciplinary Use Cases for Convergent new Approaches to AI explainability (MUCCA) project is an innovative project that aims to address this challenge by bringing together researchers from diverse fields, each contributing complementary skills essential for comprehending AI algorithm behavior. The project centers around the investigation of a wide array of multidisciplinary use-cases, where explainable AI can play a pivotal role. In our presentation we illustrate the MUCCA project in general, to then verticalize with respect to our ongoing research in the field of high energy physics. We showcase its application in both high-energy physics experiment data analysis and its use in detector and real-time systems. Specifically, we delve into our exploration of various explainability methods rooted in different approaches and evaluate their effectiveness across the diverse use-cases. The outcome of our work yields a collection of potentially comprehensible and human-friendly explanations for the predictions made by our models. We conclude by highlighting limitations of existing xAI models for high-energy physics, and brainstorming ideas on how to build novel, explainable-by-design models for accelerating scientific research with AI.

Simulation Based Inference / 6

Precision-Machine Learning for the Matrix Element Method

Auteurs: Anja Butter¹; Nathan Huetsch²; Ramon Winterhalder³; Theo Heimel²; Tilman Plehn²

¹ *LPNHE*

² *Heidelberg University*

³ *UC Louvain*

The matrix element method is the LHC inference method of choice for limited statistics. We present a dedicated machine learning framework, based on efficient phase-space integration, a learned acceptance and transfer function. It is based on a choice of INN and diffusion networks, and a transformer to solve jet combinatorics. Bayesian networks allow us to capture network uncertainties, bootstrapping allows us to estimate integration uncertainties. We showcase this setup for the CP-phase of the top Yukawa coupling in associated Higgs and single-top production.

Paper: arXiv:2310.07752

Controlling uncertainties in generative models / 7

Uncertainty-aware diffusion models for LHC Event Generation

Auteurs: Anja Butter¹; Jonas Spinner²; Nathan Huetsch²; Peter Sorrenson³; Sofia Palacios Schweitzer⁴; Tilman Plehn²

¹ *LPNHE*

² *Heidelberg University*

³ *IWR Heidelberg*

⁴ *ITP, Heidelberg University*

Given the recent success of diffusion models in image generation, we study their applicability to generating LHC phase space distributions. We find that they achieve percent level precision comparable to INNs. Training uncertainties are quantified by developing Bayesian versions to further enhance the interpretability of our results. In this talk, diffusion models are introduced and discussed followed by a presentation of our findings.

Frugal Unmixing, Reusing, and Sampling / 8

Publication and reuse of ML models in LHC analyses

Auteur: Sabine Kraml¹

¹ *LPSC Grenoble, CNRS/IN2P3*

With the increasing usage of machine learning in high energy physics analyses, the publication of the learned models in a reusable form has become a crucial question for analysis preservation and reuse. In turn, a lack of appropriate ML design and publication makes reinterpretation of analyses in terms of physics scenarios beyond those considered in the original experimental paper seriously difficult if not impossible. I will discuss recent efforts towards the preservation and reuse of ML-based LHC analyses together with guidelines for reusable ML models, which originated from the LHC Reinterpretation Forum and the 2023 PhysTeV workshop in Les Houches.

Controlling uncertainties in generative models / 9

Data driven background estimation in HEP using generative adversarial networks

Auteurs: Fabrice Couderc¹; Julie Malcles²; Mehmet Ozgur Sahin³; Victor Lohezic⁴

¹ *CEA*

² *IRFU, CEA-Saclay*

³ *Irfu, CEA Paris-Saclay*

⁴ *IRFU (CEA) / Université Paris-Saclay*

Data-driven techniques are indispensable for addressing the limitations of Monte Carlo (MC) simulations in High Energy Physics experiments, such as insufficient statistics and process mismodeling. Accurate representation of background processes is essential for achieving optimal measurement sensitivity. Traditional approaches often involve the selection of a control region to model the background, but this can introduce biases in the distribution of certain physics observables, and hence rendering them unusable in subsequent analyses. To overcome this issue, we introduced a novel method that generates physics objects that are both compatible with the region of interest and accurately represent correlations with other event properties.

To achieve this we employ conditional generative adversarial networks (GANs), leveraging their proven efficacy in various machine learning tasks. The method is illustrated by generating a new misidentified photon for the gamma+jets background of the $H \rightarrow \gamma\gamma$ analysis on the CMS Open Data simulated samples. We demonstrate that the GAN is able to generate a coherent object within the region of interest and still retains correlations with other observables within the rest of the event.

Simulation Based Inference / 10

Swyft: Direct marginal inference for large simulation models

Auteur: Christoph Weniger¹

¹ *University of Amsterdam*

As cosmology and astrophysics data advance, there is a growing demand for more detailed physical and instrumental simulation models with a multitude of uncertain parameters. Estimating the full joint posterior often becomes computationally prohibitive. Swyft is a deep learning python module that leverages the unique property of simulation-based inference to perform direct marginal inference. It enables to efficiently estimate individual parameter posteriors, perform marginal image reconstruction tasks, or do Bayesian model comparison, without access the joint posterior. I will provide a brief overview of the library and underlying algorithms, and present applications in astroparticle physics and cosmology.

Controlling uncertainties in generative models / 11

Generative modeling in genomics and a perspective on uncertainty quantification

Auteur: Burak Yelmen¹

¹ *University of Tartu*

In recent years, generative modeling has gained substantial momentum in genomics research thanks to increased availability of computational resources and development of deep generative models (DGMs) over the past decade. DGMs can learn the complex structure of genomic data and can be utilized for a variety of tasks such as generation of realistic artificial genomes, dimensionality reduction and prediction, with unsupervised, semi-supervised or supervised learning schemes. In this talk, I will present a background on generative models in genomics, discuss our recent work on the generation of artificial genomic data, and provide my perspective on approaches for uncertainty quantification.

Architectures / 12

Efficient Sampling from Bayesian Network Posteriors for Optimal Uncertainties

Auteurs: Gregor Kasieczka¹; Mathias Trabs²; Sebastian Bieringer³

¹ *Hamburg University*

² *Karlsruhe Institute of Technology*

³ *Hamburg University, Institute for experimental physics*

Bayesian neural networks are a key technique when including uncertainty predictions into neural network analysis, be it in classification, regression or generation. Although being an essential building block for classical Bayesian techniques, Markov Chain Monte Carlo methods are seldomly used to sample Bayesian neural network weight posteriors due to slow convergence rates in high dimensional parameter spaces. Metropolis-Hastings corrected chains exhibit two major issues: using a stochastic Metropolis-Hastings term and bad acceptance rates. We present solutions in form of a correction term to the loss objective and novel proposal distributions based on the Adam-optimizer. The combined algorithm shows fast convergence and good uncertainty estimation

for physics use cases without dramatically increasing the cost of computation over gradient descent based optimization.

Opening session, Uncertainty Quantification / 13

Uncertainty Quantification in Neural Networks: Methods and Considerations

Auteur: Laurens Sluiterman¹

¹ *Radboud University*

In this talk, we delve into the complexities of uncertainty quantification for neural networks. Model predictions inherently come with uncertainties that arise from several factors: stochastic outcomes, the randomness of training data samples, and the inherent variability of the training process itself. Through the lens of a regression problem, we will unpack these factors and provide a pragmatic framework to understand and quantify uncertainty effectively. Furthermore, we will discuss various considerations and pitfalls associated with using popular approaches such as Monte Carlo dropout and Deep Ensembles.

Simulation Based Inference / 14

Optimal dataset-wide inference in the presence of systematic uncertainties

Auteurs: Chris Pollard¹; Lukas Heinrich²; Philipp Windischhofer³; Siddharth Mishra-Sharma⁴

¹ *University of Warwick*

² *Technische Universität München*

³ *University of Chicago*

⁴ *Massachusetts Institute of Technology*

Real-world datasets often comprise sets of observations that collectively constrain the parameters of an underlying model of interest. Such models typically have a hierarchical structure, where “local” parameters impact individual observations and “global” parameters influence the entire dataset. In this talk we introduce Bayesian and Frequentist approaches for optimal dataset-wide probabilistic inference in cases where the likelihood is intractable, but simulations can be realized via forward modeling. We construct neural estimators for the likelihood(-ratio) or posterior and show that explicitly accounting for the model’s hierarchical structure can lead to tighter parameter constraints. We illustrate our methods using case studies from particle physics and astrophysics.

Based on: <https://arxiv.org/abs/2306.12584>

Benchmark, Datasets and Challenges / 15

Exploring Data Challenges and Leveraging Codabench: A Practical Journey with unsupervised New Physics detection at 40 MHz

Auteur: Ekaterina Govorkova¹

¹ MIT

This talk delves into our team's experience in orchestrating an unsupervised New Physics detection at 40 MHz, shedding light on the intricacies of design, implementation, and lessons learned.

We challenged the community to develop algorithms for detecting New Physics by reformulating the problem as an out-of-distribution detection task.

We provided datasets with four-vectors of the highest-momentum jets, electrons, and muons produced in a LHC collision event, together with the missing transverse energy, while the goal was to find a-priori unknown and rare New Physics hidden in a data sample dominated by ordinary Standard Model processes, using anomaly detection approaches.

We share insights gained from the past, highlighting the challenges faced and the innovative solutions employed to foster engaging and impactful competitions.

Furthermore, the presentation shifts focus to our recent exploration of Codabench as a versatile platform for orchestrating data challenges.

We share our firsthand experiences with Codabench, emphasizing its capabilities in simplifying the challenge setup process, fostering collaboration among participants, and streamlining the evaluation workflow.

Benchmark, Datasets and Challenges / 16

Neuroscience ML challenges using CodaBench: Decoding multi-limb trajectories from two-photon calcium imaging

Auteurs: Maria Dadarlat¹; Megan Lipton¹; Seungbin Park¹; Yuan-Tang Chou²

¹ Purdue University

² University of Washington, Seattle (US)

In this talk, we present the ML challenge using CodaBench for the Neuroscience dataset. The ML challenge is hosted by the Accelerated AI Algorithms for Data-Driven Discovery (A3D3) Institute.

Neural decoding is the process of predicting behavior from brain signals, which is crucial for gaining insights into the functions of various brain regions and for advancing technology, such as the brain-computer interface, to aid individuals suffering from neurological injuries and diseases. Two-photon calcium imaging is a promising technique for neural decoding to record the activity of thousands of neurons in a single-cell resolution. Dadarlatlab at Purdue University organized the dataset of neural signals recorded by two-photon calcium imaging and running trajectories of mice. Decoding two-photon calcium imaging data has been challenging because the calcium signal indirectly and non-linearly represents action potential, has slow kinematics from a long decay time of calcium fluorescence, and has low sampling rates during imaging compared to natural behavior. Artificial intelligence can be a promising solution to overcome these challenges.

17

A Data-centric AI based diagnosis method using unsupervised labeling and MLP

Auteur: mohsen zargarani¹

¹ CentraleSupelec, Institute GEEPs

The quality of Input data has a profound impact in fault diagnosis. To construct high quality datasets for required diagnosis performance, it is vital to eliminate ambiguities in featured data. An example in our case study is the undesirable similarity between I-V curves in both open-circuit fault and partial shading fault. In this study, data-centric AI architecture is shaped on both I-V curve training data

development and fault data maintenance. The diagnosis is based on an automated cluster-then-label algorithm. In the training data development step, all featured data are clustered to study the hidden meaning of each cluster representative. After assigning an initial label to each cluster, a data quality valuation runs and the assessment verifies representatives to detect qualified data with distinct fault labels. According to the assessment results, the iterative data maintenance process performs the data quality improvement. The objective of quality improvement is to eliminate the cluster ambiguities in a more smart and concise way. Our preliminary results of our data-centric AI-based diagnosis illustrate higher understandability than other state of art unsupervised fault diagnosis.

Frugal Unmixing, Reusing, and Sampling / 18

Data frugal machine learning approaches for unmixing problems in Physics

Auteur: Jerome Bobin¹

¹ CEA

Unmixing problems are ubiquitous in Physics, ranging from spectral unmixing to unsupervised component separation, where elementary physical components need to be separated out from intricate observed mixtures. These problems are generally ill-posed, which mandates the design of effective regularisation to better distinguish between the sought-after components. While ML-based methods are promising, their application is very often limited to the scarcity of the training samples (e.g. very few observations, very high cost of physical simulations, etc.). We first propose using a special type of autoencoder (AE), coined interpolatory AE, to learn adapted representations for the components to be retrieved, from very few training samples. We show how such representations can be plugged into traditional solvers to tackle unmixing problems. This will be illustrated with applications in X-ray astrophysics and spectrometry in nuclear Physics.

Controlling uncertainties in generative models / 19

Potential and challenges of highly dimensional generative models

Auteur: Gregor Kasieczka¹

¹ Universität Hamburg

Recent progress in computer science, specifically the development normalising flows and diffusion models, has brought about a breakthrough in the fidelity of generative models in particle physics.

In this talk I will first review some of these new approaches and then discuss potential uses, considering the overall theme of uncertainties. This will allow us to discuss statistical properties, performance metrics, inherent uncertainties, surrogate models and in-situ background estimation.

Architectures / 20

Robustness to Uncertainties in Machine Learning Applications for HEP

Auteur: tommaso dorigo¹

¹ INFN Sezione di Padova

In this presentation I will discuss recent trends in the handling of systematic uncertainties in HEP analysis tasks, and techniques proposed to mitigate or remove their effect in the search for optimal selection criteria and variable transformations.

The approaches discussed include nuisance-parametrized models, modified adversary losses, semi-supervised learning approaches, inference-aware techniques, and other recent developments.

Controlling uncertainties in generative models / 21

Machine-learning and equations-informed tools for generation and augmentation of turbulent data.

Auteur: Luca Biferale¹

¹ U Roma Tor Vergata

Our ability to collect data is rapidly increasing thanks to computational power and the unprecedented diversity of sensors. But how good are we at extracting, reconstructing, and understanding information from them? We present a short overview of some recent advancements for data-assimilation and modelling of turbulent multi-scale flows using both data-driven and equations-informed tools, starting from sparse and heterogeneous observations of complex fluid systems. Issues connected to validations and benchmarks in the presence of full or partial observability will be discussed. A few examples of data-generation and data-augmentation based on Generative Adversarial Learning, Diffusion Models and Nudging, for Eulerian and Lagrangian turbulence will be quantitatively discussed.

-Synthetic lagrangian turbulence by generative diffusion models

T Li, L Biferale, F Bonaccorso, MA Scarpolini, M Buzicotti
arXiv preprint arXiv:2307.08529 (2023).

- Multi-scale reconstruction of turbulent rotating flows with proper orthogonal decomposition and generative adversarial networks

T Li, M Buzicotti, L Biferale, F Bonaccorso, S Chen, M Wan
Journal of Fluid Mechanics 971, A3 (2023)

-Synchronization to big data: Nudging the Navier-Stokes equations for data assimilation of turbulent flows

PC Di Leoni, A Mazzino, L Biferale Physical Review X 10 (1), 011023 (2020)

Opening session, Uncertainty Quantification / 22

Fair Universe Challenge Paris 2023

Auteurs: Aishik Ghosh¹; Benjamin Nachman²; Benjamin Thorne^{None}; Chris Harris³; Daniel Whiteson¹; David Rousseau⁴; Elham Khoda⁵; Ihsan Ullah⁶; Isabelle Guyon⁷; Paolo Calafiura²; Peter Nugent^{None}; Ragansu Chakkappai^{None}; Sascha Diefenbacher²; Shih-Chieh Hsu⁵; Steven Farrell³; Wahid Bhimji²; Yuan-Tang Chou^{None}

¹ UC Irvine

² Lawrence Berkeley National Laboratory

³ NERSC, Berkeley Lab

⁴ IJCLab, CNRS/IN2P3, Université Paris-Saclay

⁵ University of Washington

⁶ ChaLearn

⁷ ChaLearn/Google

The Fair Universe project is building a large-compute-scale AI ecosystem for sharing datasets, training large models and hosting challenges and benchmarks. Furthermore, the project is exploiting this ecosystem for an AI challenge series focused on minimizing the effects of systematic uncertainties in High-Energy Physics (HEP), and on predicting accurate confidence intervals. This talk will describe the challenge platform we have developed that builds on the open-source benchmark ecosystem Codabench to interface it to the NERSC HPC center and its Perlmutter system with over 7000 A100 GPUs. This presentation will also mark the launch of the first of our Fair Universe public challenges hosted on this platform, the Fair Universe: HiggsML Uncertainty Challenge. There will be a hackathon during the workshop to develop the current prototype challenge, the full version of which will run in 2024. The Codabench/NERSC platform allows for hosting challenges also from other communities, and we also intend to make our benchmark designs available as templates so similar efforts can be easily launched in other domains.

Architectures / 23

Revisiting models and uncertainty with AI

Auteur: Gael Varoquaux¹

¹ *INRIA-Saclay*

Predictions from empirical evidence come with many sources of potential uncertainty and error. First, the specific choices of models and concepts that we tack onto the observation give a strong prism to the resulting conclusion. Uncertainty on which functional form to use in a model, naturally results in uncertainty of conclusions. Outside of mature (post-paradigmatic) quantitative sciences such as physics, the mere choice of ingredients put the model (which quantities to measure) is open.

I will discuss how AI, or machine learning, brings a new angle to these questions, because it tackles complex observations with very flexible models. I believe that it opens new doors to scientific evidence by putting the burden on validity on model outputs, rather than ingredients.

However, a given model fitted on data should ideally express its uncertainty as a probability of the output given the input. This is particularly important in high-stakes applications such as health. I would discuss how controlling this uncertainty requires to control a quantity known as calibration, but also to go further and control the remainder, the “grouping loss”, which leads to challenging estimation problems.

Architectures / 24

Data Subsampling for Bayesian Neural Networks

Auteur: Eiji Kawasaki¹

Co-auteur: Markus Holzmann²

¹ *CEA*

² *Univ. Grenoble Alpes, CNRS, LPMMC, 38000 Grenoble, France*

The development of an effective Uncertainty Quantification method that computes the predictive distribution by marginalizing over Deep Neural Network parameter sets remains an important, challenging task. In this context, Markov Chain Monte Carlo algorithms do not scale well for large datasets leading to difficulties in Neural Network posterior sampling. During this talk, we’ll show that a generalization of the Metropolis Hastings algorithm allows to restrict the evaluation of the likelihood to small mini-batches in a Bayesian inference context. Since it requires the computation of a so-called “noise penalty” determined by the variance of the training loss function over the

mini-batches, we refer to this data subsampling strategy as Penalty Bayesian Neural Networks – PBNNs.

Closing session / 25

Fair Universe HiggsML Uncertainty Challenge: Lessons Learned and plans

Auteurs: Benjamin Nachman¹; David Rousseau²; Elham E Khoda³; Ragansu Chakkappai^{None}; Sascha Diefenbacher¹; Wahid Bhimji¹

¹ *Lawrence Berkeley National Laboratory*

² *IJCLab, CNRS/IN2P3, Université Paris-Saclay*

³ *University of Washington*

The first FAIR Universe public challenge, FAIR Universe: HiggsML Uncertainty Challenge, will be launched before this workshop. A dedicated hackathon within the workshop will refine the existing prototype, with the full version set for release in 2024. The intricacies of devising an efficient scoring method for uncertainty-aware techniques pose a challenge, and the current approach will be discussed during the workshop. This presentation will spotlight insights gained from diverse workshop activities and explore potential approaches for forthcoming iterations of the FAIR Universe challenge.

Fair-Universe hackathon / 26

Overview of HiggsML Uncertainty Challenge

This afternoon's hackathon will center around the prototype "HiggsML Uncertainty Challenge" to be fully launched in 2024. This overview will describe the setup of the prototype challenge.

Opening session, Uncertainty Quantification / 27

Uncertainty Quantification in High Energy Physics

Opening session, Uncertainty Quantification / 28

Bayesian optimisation

Opening session, Uncertainty Quantification / 29

Uncertainty Quantification in Industry

Opening session, Uncertainty Quantification / 30**Welcome****Simulation Based Inference / 31****Simulation-Based Inference: Where Classical Statistics Meets Machine Learning****Auteur:** Mikael Kuusela¹¹ *Carnegie Mellon University*

Simulation-based inference (SBI) refers to situations where the likelihood function cannot be readily evaluated but a simulator is available to generate data from a parametric model for any value of the unknown parameter. In recent years, a wide range of machine learning-based techniques have been developed to enable classical statistical inference in the simulation-based setting. These methods enable learning likelihood functions, posteriors, likelihood ratios, confidence sets and other inferential quantities with very high-dimensional data spaces using only simulations from the statistical model. In this talk, I will give an overview of some of our recent work in this area. I will first present our work on learning likelihood functions of otherwise intractable models for spatial data, which is one of the first uses of SBI in the context of purely statistical models. While likelihood-based inference is conceptually simple and computationally efficient, it only offers approximate coverage guarantees. To obtain rigorous uncertainty quantification, one can combine SBI with Neyman inversion. From this class of techniques, I will present Waldo, a method that uses neural predictions to form a Wald-type test statistic which is inverted to obtain guaranteed-coverage confidence sets. Waldo provides an appealing way of obtaining frequentist uncertainty quantification based on outputs of predictive models, including deep neural networks and neural posteriors. I will conclude by outlining some future research directions in SBI which revolve around challenges related to model discrepancy and high-dimensional parameter spaces.

Explainable AI / 32**Explainable AI for Interpretability of Deep Neural Networks : the High Energy Physics perspective**

Explainable AI (xAI) represents a set of processes and methods that allows human users to comprehend results created by machine learning algorithms. In the context of applications of AI to science, we need to look beyond standard metrics of prediction performance such as accuracy to ensure that AI models are robust to noise and adversarial samples, fair to biases in data populations, and generate trustworthy explanations of their predictions. A challenge is that xAI is hard to define and even harder to evaluate. There is no universal definition of what it means for an AI model to be explainable nor well-defined metrics to evaluate the “goodness” of explanations generated for AI models. Despite of these challenges, current xAI tools and methods are powerful allies for physicists. They have great utility in aiding in the interpretation deep neural networks (DNNs) and this information can be used to create better algorithms.

In this talk, I will discuss these aspects of xAI and the application of xAI methods to DNN models used in jet tagging. In our case study of jets coming from top quark decay in the high energy proton-proton collisions at the Large Hadron Collider, we use XAI to help identify which features play the most important roles in identifying the top jets, how and why feature importance varies across different XAI metrics, and how latent space representations encode information as well as correlate with physical quantities. We additionally illustrate the activity of hidden layers as Neural Activation Pattern (NAP) diagrams to understand how DNNs relay information across the layers and how this

understanding can help us to make such models significantly simpler by allowing effective model re-optimization and hyperparameter tuning.

Explainable AI / 33

xAI in practice: current state of the art, limitations and perspectives

Explainable Artificial Intelligence (xAI) is a vibrant research field that aims to provide an insight on the decision taken by machine learning (ML) programs. The pervasiveness of AI in our societies pushed regulations (as the European AI Act) that demands transparency. We will present an overview of the field of Explainable AI, focusing on local explanations. We will present the various caveat identified by the literature, and some of our perspectives on xAI.

Simulation Based Inference / 34

SBI

Simulation Based Inference / 35

Simulation Based Inference: The Frequentist Perspective Abstract:

I give a brief introduction to the frequentist approach to simulation-based inference, which is often referred to as likelihood-free frequentist inference. The approach is illustrated with three simple examples, one from cosmology, one from particle physics, and the third from epidemiology.

Simulation Based Inference / 36

SBI Introduction

Frugal Unmixing, Reusing, and Sampling / 37

TBC

Unfolding (de-biasing) / 38

Unfolding in High Energy Physics

In high-energy physics, unfolding is a critical statistical process for interpreting experimental data that is complicated by the intrinsic ill-posedness of the problem. This complexity arises from the need to provide heuristics for statistical estimates that disentangle true physical phenomena from observational distortions. We present a typical roadmap for why, when, and how unfolding is applied in high energy physics experiments and how the treatment of uncertainties influences considerations such as the choice of algorithm and regularisation. Finally, the concept of unbinned unfolding is presented together with a description of how statistical and systematic uncertainties are typically addressed in unfolding problems, together with a discussion of how statistical modelling and AI can lead to better estimates in the future.

Unfolding (de-biasing) / 39

Uncertainty Quantification and Anomaly Detection with Evidential Deep Learning

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 will discuss its connection with anomaly detection (AD). Several examples will be presented, including ongoing work in this area for HEP applications.

Closing session / 40

Highlights from the 1st CNRS AISSAI Thematic Quarter on Causality

Causal and effect questions are the cornerstone of numerous scientific disciplines, providing a framework for formulating and comprehending them under diverse conditions. In recent years, a notable interdisciplinary effort has been to develop new methods to address causality-related challenges. The first CNRS AISSAI Thematic Quarter on Causality, held earlier this year, marked a significant interdisciplinary step in advancing our understanding of causal relationships. This talk aims to highlight some insights and methodologies gleaned during the quarter. They include the use of sigma-algebra via the Witsenhausen Intrinsic Model (WIM) and the Information Dependency Model (IDM) as an alternative approach to the traditional functional causal models and causal graphs, the axiomatization of causality through Kolmogorov's measure-theoretic axiomatization of probability, and the use of cumulant tensors for characterizing hidden common causes when learning causal graphs in linear non-Gaussian causal models. Moreover, the quarter underscored the promising role of diffusion and normalizing flow methods in capturing the underlying causal data-generating processes, offering new horizons in causal representation learning.

Frugal Unmixing, Reusing, and Sampling / 41

Machine Learning Assisted Sampling: Applications in Physics

This workshop gathered experts in computational methods for rare events sampling interested in machine learning methods and experts in generative models interested in applications to physical

systems. In this talk I will try to summarize the main directions of research that were highlighted in the workshops: how learning can help for dimensionality reduction? which machine learning strategies can be designed to sample rare events? and how generative models can be adapted to physics-applications?

Closing session / 42

Fair Universe hackathon : outcome and plans

Closing session / 43

Farewell

Opening session, Uncertainty Quantification / 44

Uncertainty modeling in particle physics

Auteur: Wouter Verkerke¹

¹ *Nikhef/UvA*

I will present a pedagogical introduction to uncertainty modeling in particle physics. I will mostly focus on the methods used at the Large Hadron Collider experiments, where systematic effects are explicitly parameterized in the likelihood function in terms of nuisance parameters. Accurate modeling of systematic effects is of increasing importance at the LHC as the abundant data has decreased statistical uncertainties in many measurements to be on par with systematic uncertainties. I will discuss the reasoning behind the modeling approaches commonly chosen, common challenges in the parametric modeling and in the interpretation of the corresponding uncertainties. I will conclude with the special considerations in the modeling of theoretical uncertainties, which are often incompletely defined.

Opening session, Uncertainty Quantification / 46

An Introduction to Bayesian Optimization

Auteur: Emmanuel Vazquez¹

¹ *L2S – Paris-Saclay*

In this talk, we delve into the foundational principles of Bayesian optimization, a method particularly well-suited for optimizing deterministic or stochastic functions, whether scalar or vectorial, especially when the evaluation of the function is computationally expensive and no gradient information is available.

Bayesian optimization is particularly relevant in the domains of Design and Analysis of Computer Experiments (DACE) and Uncertainty Quantification (UQ). In these areas, it is typically applied to minimize costs or maximize performance through complex computer simulations, including those

utilizing Partial Differential Equations and finite element methods. Additionally, Bayesian optimization has become increasingly popular in Machine Learning and Artificial Intelligence for optimizing parameters in learning procedures, highlighting its versatility and effectiveness across a range of fields.

At the core of Bayesian optimization is the practice of modeling the target function using the Gaussian process framework. This modeling approach enables the construction of a sampling criterion, also known as an acquisition function. We will examine various classic sampling criteria, including the widely-used Expected Improvement (EI) criterion, and discuss the concept of Stepwise Uncertainty Reduction (SUR). Our discussion will focus on how Bayesian optimization effectively manages the balance between exploration and exploitation.

Fair-Universe hackathon / 47

Walkthrough of CodaBench and submissions

Join us for an interactive workshop featuring a hands-on hackathon and insightful discussions. The session kicks off with a comprehensive tutorial on Codabench, an open-source challenge organization platform (codabench.org). During this tutorial, participants will gain familiarity with various aspects of the platform, including:

- * Signing up/signing in to Codabench
- * Navigating the overall structure of the platform
- * Following the step-by-step getting started instructions
- * Accessing sample submissions
- * Utilizing the provided starting kit
- * Submitting pre-existing submissions
- * Exploring the leaderboard for scores

Once the hands-on tutorial concludes, participants are encouraged to apply their newfound skills by engaging in the HiggsML Uncertainty Challenge. Organizers will be onsite to provide assistance and guidance, ensuring a collaborative and enriching experience for all attendees.

Fair-Universe hackathon / 48

Scoring and Baseline Systematic Aware method

Fair-Universe hackathon / 49

Feedback from participants

Fair-Universe hackathon / 50

Hands-on hackathon and discussion

Fair-Universe hackathon / 51

Wrap up and next steps

Opening session, Uncertainty Quantification / 52

Uncertainty Quantification and Machine Learning in Industry: Current Practices and Challenges in Industrial Applications for Low-Carbon Electricity Production

Auteur: Vincent CHABRIDON¹

Co-auteurs: Antoine Ajenjo ¹; Michaël Baudin ¹; Nicolas Bousquet ¹; Elias Fekhari ¹; Bertrand Iooss ¹; Merlin Keller ¹; Joseph Muré ¹; Julien Pelamatti ¹; Emmanuel Remy ¹; Roman Sueur ¹

¹ EDF R&D

This talk will focus on a panel of current practices and challenges regarding both Uncertainty Quantification (UQ) and Artificial Intelligence (mainly from the Machine Learning (ML) point of view), in EDF's industrial applications, especially in the topics of risk management of industrial production assets. From our point of view, these two core topics, UQ & AI are, today, closely related to one another, especially when targeted applications are critical industrial systems (such as, e.g., nuclear reactors, dams, or wind farms). After presenting the commonly accepted methodology for UQ (and its link with ML) in numerical simulation codes, this talk will try to provide a panel of motivating applications that EDF R&D is confronted with, together with the industrial problems they arise. Then, a variety of technical and scientific challenges will be derived from these applications to introduce several research tracks we developed and pursued in the last decade. In addition to these methodological contributions, an emphasis will be put on the various open-source tools and software that have been produced as byproducts of this long-term research endeavour. Finally, a few current open questions will be discussed at the end of the talk while opening the path to future research questions and applications.

Simulation Based Inference / 53

An introduction to simulation-based inference

Auteur: Gilles Louppe¹

¹ University of Liège

In this talk, we will introduce simulation-based inference and present how deep learning can be used to solve complex inverse problems commonly found in scientific disciplines. We will give an introduction and overview of the topic and present some of our recent work on the topic. We will also discuss the opportunities and challenges.

Frugal Unmixing, Reusing, and Sampling / 54

Introduction at Institut Pascal

Fair-Universe hackathon / 55

Introduction