



**Uncertainty Quantification
for
Machine Learning**

UQ4ML



CEA Paris-Saclay, France
Institut de Physique Théorique (IPhT)
Amphithéâtre Claude Bloch
15 – 18 September 2025



Contents

Simulation and Coding	5
Quantifying uncertainties in computer models: an overview	6
Uncertainty in AI-driven Monte Carlo simulations	7
Parameter Estimation with Neural Simulation-Based Inference in ATLAS	8
HEP – Theory	9
Machine learning for optimal BSM sensitivity	10
Determination and validation of modeling and theory uncertainties	11
Parton Distributions from Neural Networks: Analytical Results	12
Energy Flow Networks for Jet Quenching Studies	13
HEP – Experiment	15
Uncertainty quantification for deep learning in astroparticle physics	16
Interdisciplinary Digital Twin Engine InterTwin for calorimeter simulation	17
Amplitude Learning Cookbook	18
Fair Universe HiggsML Uncertainty Challenge	19
Data Analysis, Time series, Causal analysis	21
Causal Discovery for Time Series Based on HSIC and Kernel Approximations	22
Calibrated and uncertain? Evaluating uncertainty estimates in binary classification models	23
Unsupervised Anomaly Detection in Multivariate Time Series Using Public Benchmarks and Synthetic Data from Lorenzetti	24
Bayesian continual learning and forgetting in neural networks	25
Deep Learning and Uncertainty Quantification	27
Uncertainty Quantification for Neural Networks in Particle Physics	28
Uncertainty Quantification in an ML Pattern Recognition Pipeline	29
Uncertainty quantification for machine learning: a new approach for the Critical Heat Flux application	30
Spotlight Talks	31
Machine learning of biological sequences	32
Interaction reconstruction in scintillator detectors for PET imaging – Deep Learning approach with uncertainty quantification	33

Simulation and Coding

Monday, 15 September 2025, 14:15 – 16:45

Chair: **Dr. Jean-Baptiste Blanchard** (CEA Paris-Saclay)

This session presents an overview of uncertainty quantification techniques for simulations in machine learning. By examining how uncertainties propagate through complex models, the session explores the reliability and limitations of their results. This knowledge is crucial for advancing the field and ensuring that machine learning applications are robust and trustworthy.

Quantifying uncertainties in computer models: an overview

Dr. Merlin Keller (EDF Paris)

Keynote

Monday, 15 September 2025, 14:15

Predictions from computer models are now extensively used in industrial studies, to complement, or even sometimes replace, field experiments. Such numerical experiments have key advantages, such as reduced costs, and added flexibility. However, they raise the question of assessing the validity of computer model predictions, with respect to the physical phenomena they seek to reproduce. This is the goal of verification, validation and uncertainty quantification (VVUQ), a process whose development is an active field of research in applied mathematics. The goal of this talk is to present the different objectives and challenges of VVUQ, as well as a generic workflow that can be applied to virtually any uncertainty quantification problem. We then discuss software libraries that implement this methodology, as well as open problems and perspectives.

Uncertainty in AI-driven Monte Carlo simulations

Dimitrios Tzivraillis (CEA Paris-Saclay)

Alberto Rosso (CNRS, Université Paris-Saclay)

Eiji Kawasaki (CEA Paris-Saclay)

Short-talk

Monday, 15 September 2025, 15:45

In the study of complex systems, evaluating physical observables often requires sampling representative configurations via Monte Carlo techniques. These methods rely on repeated evaluations of the system's energy and force fields, which can become computationally expensive. To accelerate these simulations, deep learning models are increasingly employed as surrogate functions to approximate the energy landscape or force fields. However, such models introduce epistemic uncertainty in their predictions, which may propagate through the sampling process and affect the system's macroscopic behavior. In our work, we present the Penalty Ensemble Method (PEM) to quantify epistemic uncertainty and mitigate its impact on Monte Carlo sampling. Our approach introduces an uncertainty-aware modification of the Metropolis acceptance rule, which increases the rejection probability in regions of high uncertainty, thereby enhancing the reliability of the simulation outcomes. Abstract to be defined.

Parameter Estimation with Neural Simulation-Based Inference in ATLAS

David Rousseau (IJCLab-Orsay)

Short-talk

Monday, 15 September 2025, 16:15

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.

This talk covers [arXiv:2412.01600](https://arxiv.org/abs/2412.01600) which is a methodology paper detailing the NSBI technique used in the physics paper [arXiv:2412.01548](https://arxiv.org/abs/2412.01548).

HEP – Theory

Tuesday, 16 September 2025, 9:30 – 12:30

Chair: **Dr. Marco Letizia** (University of Genoa)

This session explores the theoretical aspects of high energy physics through the lens of machine learning. By addressing the unique challenges and uncertainties in this domain, we review innovative approaches to enhance our understanding of fundamental physics. This intersection of fields highlights the broader implications of uncertainty quantification in scientific discovery.

Machine learning for optimal BSM sensitivity

Prof. Giuliano Panico (University of Florence and INFN Florence)

Keynote

Tuesday, 16 September 2025, 9:30

I will review the parametrized classifiers for optimizing the sensitivity to EFT operators and some the machine-learning approaches for general anomaly detection. Particular attention will be devoted to validation procedures and ways to treat uncertainties.

Determination and validation of modeling and theory uncertainties

Prof. Stefano Forte (Università degli Studi e INFN Milano)

Short-talk

Tuesday, 16 September 2025, 11:00

I discuss how uncertainties related to machine learning modeling of a regression problem, as well as those related to missing theoretical information, can be estimated and subsequently validated. Even though these uncertainties are intrinsically Bayesian, given that there is only one underlying true theory and true model, they can be determined both in a Bayesian and frequentist framework. I show how this can be done in the context of the determination of the parton distributions that encode the structure of the proton. I further show how results can be validated by means of closure tests.

Parton Distributions from Neural Networks: Analytical Results

Amedeo Chiefa (The University of Edinburgh)

Prof. Luigi Del Debbio (The University of Edinburgh)

Prof. Richard Kenway (The University of Edinburgh)

Short-talk

Tuesday, 16 September 2025, 11:30

Parton Distribution Functions (PDFs) play a crucial role in describing experimental data at hadron colliders and provide insight into proton structure. As the LHC enters an era of high-precision measurements, a robust PDF determination with a reliable uncertainty quantification has become increasingly important to match the experimental precision. The NNPDF collaboration has pioneered the use of Machine Learning (ML) techniques for PDF determination. In this work, we develop a theoretical framework based on the Neural Tangent Kernel (NTK) to analyse the training dynamics of Neural Networks. This approach allows us to derive, under certain assumptions, an analytical description of how the neural network evolves during training, enabling us to better understand the NNPDF methodology and its dependence on the underlying model architecture. Notably, we demonstrate that our results contrast, to some extent, with the standard picture of the lazy training regime commonly discussed in the ML community.

Energy Flow Networks for Jet Quenching Studies

João A. Gonçalves (LIP - IST)

Short-talk

Tuesday, 16 September 2025, 12:00

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 the primordial liquid. In this study, we 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 a crucial step towards separating HI jets that were quenched from those with little or no modification by the interaction with the QGP on a jet-by-jet basis. We trained simple Energy Flow Networks (EFNs) and further enhanced them by incorporating jet observables such as N-Subjettiness and Energy Flow Polynomials (EFPs). Our primary objective is to assess the effectiveness of these approaches in the context of Jet Quenching, exploring new phenomenological avenues by combining these models with various encodings of jet information. Initial evaluations using Linear Discriminant Analysis (LDA) set a performance baseline, which is significantly enhanced through simple Deep Neural Networks (DNNs), capable of capturing non-linear relations expected in the data. Integrating both 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. This significant performance is noteworthy given that both medium response and underlying event contamination effects on the jet are taken into account. These results underscore the potential of combining EFNs with jet substructure observables to advance Jet Quenching studies and adjacent areas, paving the way for deeper insights into the properties of the QGP. Results on a variation of EFNs, Moment EFNs (MEFNs), which can achieve comparable performance with a more manageable and, in turn interpretable, latent space, will be presented.

HEP – Experiment

Tuesday, 16 September 2025, 14:00 – 17:00

Chair: **Dr. Shamik Ghosh** (Centre National de la Recherche Scientifique (CNRS))

This session delves into the practical challenges and issues of applying AI in experimental high energy physics. Seminars will deal with how uncertainty quantification plays a pivotal role in ensuring the accuracy and reliability of experimental results. This discussion is essential for advancing both the field of high energy physics and the broader application of machine learning in experimental sciences.

Uncertainty quantification for deep learning in astroparticle physics

Prof. Christian Glaser (Uppsala University)

Keynote

Tuesday, 16 September 2025, 14:00

In this contribution I will review the use cases of uncertainty quantification with deep learning in high-energy astroparticle physics. Among other things, I will present the combination of neural networks with conditional normalizing flows to predict the Posterior for all quantities of interest. This Ansatz can be further expanded with the snowstorm method developed by the IceCube collaboration to include systematic uncertainties in the Posterior prediction by sampling from the systematic uncertainties during MC generation.

Interdisciplinary Digital Twin Engine InterTwin for calorimeter simulation

Vera Maiboroda (CNRS, IJCLab)

Corentin Allaire (IJCLab, Université Paris-Saclay, CNRS/IN2P3)

David Rousseau (IJCLab-Orsay)

Short-talk

Tuesday, 16 September 2025, 15:30

The interTwin project develops an open-source Digital Twin Engine to integrate application-specific Digital Twins (DTs) across scientific domains. Its framework for the development of DTs supports interoperability, performance, portability and accuracy. As part of this initiative, we implemented the CaloINN normalizing-flow model for calorimeter simulations within the interTwin framework. Calorimeter shower simulations are computationally expensive, and generative models offer an efficient alternative. However, achieving a balance between accuracy and speed remains a challenge, with distribution tail modeling being a key limitation. CaloINN provides a trade-off between simulation quality and efficiency. The ongoing study targets validating the model using high granularity simulations from the Open Data Detector, as well as introducing a set of post-processing modifications of analysis-level observables aimed at improving the accuracy of distribution tails.

Amplitude Learning Cookbook

Joaquin Iturriza Ramirez (CNRS)

Prof. Anja Butter (CNRS)

Prof. Bertrand Laforge (LPNHE - Sorbonne Université)

Prof. Henning Bahl (ITP Heidelberg)

Prof. Tilman Plehn (ITP Heidelberg)

Dr. Víctor Bresó Pla (ITP Heidelberg)

Short-talk

Tuesday, 16 September 2025, 16:00

Fast and precise evaluations of scattering amplitudes even in the case of precision calculations is essential for event generation tools at the HL-LHC. We explore the scaling behavior of the achievable precision of neural networks in this regression problem for multiple architectures, including a Lorentz symmetry aware multilayer perceptron and the L-GATr architecture. L-GATr is equivariant with respect to the Lorentz group by its internal embedding living in the geometric algebra defined by the flat space-time metric. This study addresses in particular the scaling behavior of uncertainty estimations using state of the art methods.

Fair Universe HiggsML Uncertainty Challenge

Ragansu Chakkappai (IJCLab-Orsay)

Prof. Aishik Ghosh (University of California Irvine)

Dr. Ben Nachman (Lawrence Berkeley National Laboratory)

Dr. Chris Harris (NERSC)

Prof. David Rousseau (IJCLab-Orsay)

Dr. Elham Khoda (University of Washington)

Dr. Ihsan Ullah (Chalearn)

Dr. Isabelle Guyon (Chalearn)

Jordan Dudley (Lawrence Berkeley National Laboratory)

Dr. Paolo Calafiura (Lawrence Berkeley National Laboratory)

Dr. Peter Nugent (Lawrence Berkeley National Laboratory)

Dr. Po-Wen Chang (Lawrence Berkeley National Laboratory)

Dr. Sascha Diefenbacher (Lawrence Berkeley National Laboratory)

Prof. Shih-Chieh Hsu (University of Washington Seattle)

Dr. Steven Farrell (Lawrence Berkeley National Laboratory)

Prof. Wahid Bhimji (Lawrence Berkeley National Laboratory)

Dr. Yuan-Tang Chou (University of Washington)

Dr. Yulei Zhang (University of Washington)

Short-talk

Tuesday, 16 September 2025, 16:30

The Fair Universe project organised the HiggsML Uncertainty Challenge, which took place from 12th September 2024, to 14th March 2025. This groundbreaking competition in high-energy physics (HEP) and machine learning was the first to strongly emphasize on uncertainties, focusing on mastering both the uncertainties in the input training data and providing credible confidence intervals in the results. The challenge revolved around measuring the Higgs to $\tau_+ \tau_-$ cross-section, similar to the HiggsML challenge held on Kaggle in 2014, using a dataset representing the 4-momentum signal state. Participants were tasked with developing advanced analysis techniques capable of measuring the signal strength and generating confidence intervals that included both statistical and systematic uncertainties, such as those related to detector calibration and background levels. The accuracy of these intervals was automatically evaluated using pseudo-experiments to assess correct coverage. Techniques that effectively managed the impact of systematic uncertainties were expected to perform best, contributing to the development of uncertainty-aware AI techniques for HEP and potentially other fields. The competition was hosted on Codabench, an evolution of the Codalab platform, and leveraged significant resources from the NERSC infrastructure to handle the thousands of required pseudo-experiments. This competition was selected as a NeurIPS competition, and the preliminary results were presented at the NeurIPS 2024 conference in December. As the challenge concluded in March 2025, an account of the winning solutions will be presented at this workshop.

Data Analysis, Time series, Causal analysis

Wednesday, 17 September 2025, 9:30 – 12:30

Chair: **Dr. Riccardo Finotello** (CEA Paris-Saclay)

This session focuses on the critical aspects of data analysis in time series, though it is not limited to it. Understanding and quantifying uncertainty in time-dependent data is essential for accurate forecasting and decision-making. Additionally, exploring causal relationships helps in identifying true drivers of outcomes, rather than mere correlations, which is vital for robust machine learning models.

Causal Discovery for Time Series Based on HSIC and Kernel Approximations

Dr. Aurore Lomet (CEA Paris-Saclay)

Keynote

Wednesday, 17 September 2025, 9:30

Causality, in Pearl's framework, is defined through structural causal models: systems of structural equations with exogenous variables and a directed acyclic graph that encodes cause-effect relations. In contrast, correlation, which often forms the basis of artificial intelligence models, quantifies statistical association and may arise from confounding or indirect paths without implying a directed effect. In physics, causal analysis enables the identification of mechanisms that support predictive models under interventions. Correlation alone cannot determine whether an observed dependency originates from a physical mechanism, a hidden common cause, or a statistical artifact. Moreover, consideration of causal structure provides a basis for interpretable artificial intelligence models, since the reasoning aligns with the formal description of interactions in physics. In this context, causal discovery aims to recover the underlying causal graph from data. Approaches based on independence testing require measures that remain valid under realistic assumptions. Linear methods such as partial correlation impose Gaussianity and linear relations. Kernel-based measures such as the Hilbert-Schmidt Independence Criterion (HSIC) relax these constraints and detect nonlinear dependencies. Building on this foundation, this presentation introduces a causal discovery method for time series, such as sensor data, that employs kernel dependence measures, integrates kernel approximations to improve computational efficiency, and takes into account the temporal dependence in the HSIC to reduce bias from autocorrelation. The results indicate that the proposed method achieves consistent performances across different levels of dependence and length of time series compared to alternative approaches.

Calibrated and uncertain? Evaluating uncertainty estimates in binary classification models

Aurora Singstad Grefsrud (Western Norway University of Applied Sciences)

Prof. Nello Blaser (University of Bergen)

Dr. Trygve Buanes (Western Norway University of Applied Sciences)

Short-talk

Wednesday, 17 September 2025, 11:00

Rigorous statistical methods, including the estimation of parameter values and their uncertainties, underpins the validity of scientific discovery, and has been especially important in the natural sciences. In the age of data-driven modeling, where the complexity of data and statistical models grow exponentially as computing power increases, uncertainty quantification has become exceedingly difficult and a plethora of techniques using a wide variety of more or less developed mathematical foundations have been proposed. In this case study we use the unifying theoretical framework of (approximate) nonparametric Bayesian inference and empirical tests on carefully created synthetic datasets to investigate qualitative properties of 6 different probabilistic machine learning algorithms for class probability and uncertainty estimation: (i) a neural network ensemble, (ii) neural network ensemble with conflictual loss, (iii) evidential deep learning, (iv) a single neural network with Monte Carlo Dropout, (v) Gaussian process classification and (vi) a Dirichlet process mixture model. We check if the algorithms produce results which reflect commonly desired statistical properties of uncertainty estimates in these kind of models such as calibration and an increase in uncertainty for out of distribution data points. Our results indicate that all algorithms are well calibrated, but none of the deep learning based algorithms provide uncertainties that reliably reflect lack of experimental evidence for out of distribution data points. We hope our study may serve as a clarifying example for researchers in the natural sciences trying to navigate the field of uncertainty quantification and especially to those developing new methods of uncertainty estimation for scientific data-driven modeling.

Unsupervised Anomaly Detection in Multivariate Time Series Using Public Benchmarks and Synthetic Data from Lorenzetti

Laura Boggia (CNRS)

Prof. Bogdan Malaescu (LPNHE-Paris CNRS/IN2P3)

Short-talk

Wednesday, 17 September 2025, 11:30

Anomaly detection in multivariate time series is an important problem across various fields such as healthcare, financial services, manufacturing or physics detector monitoring. Accurately identifying the instances when defects occur is essential but challenging, as the types of anomalies are unknown beforehand and reliably labelled data are scarce. We evaluate unsupervised transformer-based models and benchmark their performance against traditional methods on public data. Furthermore, to address the lack of reliable labels, we use the Lorenzetti Shower simulator - a general-purpose framework for simulating high-energy calorimeters - where we introduce artificial defects to evaluate the sensitivity of various detection methods.

Bayesian continual learning and forgetting in neural networks

Kellian Cottart (Université Paris Saclay)

Dr. Damien Querlioz (Université Paris Saclay)

Dr. Djohan Bonnet (Forschungszentrum Jülich)

Dr. Elisa Vianello (CEA LETI)

Tharcisius Januel (CEA LETI)

Dr. Thomas Dalgaty (CEA LETI)

Dr. Tifenn Hirtzlin (CEA LETI)

Short-talk

Wednesday, 17 September 2025, 12:00

Biological synapses effortlessly balance memory retention and flexibility, yet artificial neural networks still struggle with the extremes of catastrophic forgetting and catastrophic remembering. Here, we introduce Metaplasticity from Synaptic Uncertainty (MESU), a Bayesian framework that updates network parameters according to their uncertainty. This approach allows a principled combination of learning and forgetting that ensures that critical knowledge is preserved while unused or outdated information is gradually forgotten. Unlike standard Bayesian approaches – which risk becoming overly constrained because their posterior variances keep shrinking as evidence from all past tasks accumulates, and popular synaptic-consolidation-based continual-learning methods that rely on explicit task boundaries, MESU seamlessly adapts to streaming data. It further provides reliable epistemic uncertainty estimates, allowing out-of-distribution detection, the only computational cost being to sample the weights multiple times to provide proper output statistics. Experiments on image-classification benchmarks demonstrate that MESU mitigates catastrophic forgetting, while maintaining plasticity for new tasks. When training 200 sequential Permuted MNIST tasks, MESU outperforms established synaptic-consolidation-based continual learning techniques in terms of accuracy, capability to learn additional tasks, and out-of-distribution data detection. Additionally, due to its non-reliance on task boundaries, MESU outperforms conventional learning techniques on the incremental training of CIFAR-100 tasks consistently in a wide range of scenarios. Our results unify ideas from metaplasticity, Bayesian inference, and Hessian-based regularization, offering a biologically-inspired pathway to robust, perpetual learning.

Deep Learning and Uncertainty Quantification

Wednesday, 17 September 2025, 14:00 – 17:00

Chair: **Dr. Alessandra Cappati** (CEA Paris-Saclay)

Deep learning models have shown remarkable success in various domains, but they often lack transparency in their predictions. This session addresses the integration of uncertainty quantification techniques with deep learning, enabling more reliable and interpretable models.

Uncertainty Quantification for Neural Networks in Particle Physics

Dr. Anja Butter (CNRS)

Keynote

Wednesday, 17 September 2025, 14:00

Correctly calibrated uncertainties have always been a fundamental pillar of particle physics. As machine learning becomes increasingly integrated into both experimental and theoretical workflows, it is essential that neural network predictions include robust and reliable uncertainty estimates.

This talk will review current approaches to uncertainty estimation in neural networks, focusing on Bayesian neural networks, heteroscedastic loss functions, and repulsive ensembles. Their calibration and practical challenges will be discussed through examples from amplitude regression and unfolding. Additionally, we will explore how machine learning concepts of aleatoric and epistemic uncertainty relate to the statistical and systematic uncertainties familiar in particle physics.

Uncertainty Quantification in an ML Pattern Recognition Pipeline

Lukas Péron (ENS Paris, Lawrence Berkeley National Laboratory)

Dr. Jay Chan (Lawrence Berkeley National Laboratory)

Dr. Paolo Calafiura (Lawrence Berkeley National Laboratory)

Dr. Xiangyang Ju (Lawrence Berkeley National Laboratory)

Short-talk

Wednesday, 17 September 2025, 15:30

Geometric learning pipelines have achieved state-of-the-art performance in High-Energy and Nuclear Physics reconstruction tasks like flavor tagging and particle tracking [1]. Starting from a point cloud of detector or particle-level measurements, a graph can be built where the measurements are nodes, and where the edges represent all possible physics relationships between the nodes. Depending on the size of the resulting input graph, a filtering stage may be needed to sparsify the graph connections. A Graph Neural Network will then build a latent representation of the input graph that can be used to predict, for example, whether two nodes (measurements) belong to the same particle or to classify a node as noise. The graph may then be partitioned into particle-level subgraphs, and a regression task used to infer the particle properties. Evaluating the uncertainty of the overall pipeline is important to measure and increase the statistical significance of the final result. How do we measure the uncertainty of the predictions of a multistep pattern recognition pipeline? How do we know which step of the pipeline contributes the most to the prediction uncertainty, and how do we distinguish between irreducible uncertainties arising from the aleatoric nature of our input data (detector noise, multiple scattering, etc) and epistemic uncertainties that we could reduce by using, for example, a larger model, or more training data?

We have developed an Uncertainty Quantification process for multistep pipelines to study these questions and applied it to the acorn particle tracking pipeline [2]. All our experiments are made using the TrackML open dataset [3]. Using the Monte Carlo Dropout method, we measure the data and model uncertainties of the pipeline steps, study how they propagate down the pipeline, and how they are impacted by the training dataset's size, the input data's geometry and physical properties. We will show that for our case study, as the training dataset grows, the overall uncertainty becomes dominated by aleatoric uncertainty, indicating that we had sufficient data to train the acorn model we chose to its full potential. We show that the ACORN pipeline yields high confidence in the track reconstruction and does not suffer from the miscalibration of the GNN model.

[1] [Graph Neural Networks in Particle Physics: Implementations, Innovations, and Challenges](#)

[2] [acorn - GNN4ITkTeam](#)

[3] [Data - TrackML particle tracking challenge](#)

Uncertainty quantification for machine learning: a new approach for the Critical Heat Flux application

Michele Cazzola (CEA Paris-Saclay)

Gabriele Cassetta (CEA Paris-Saclay)

Dr. Alberto Ghione (CEA Paris Saclay)

Dr. Julien Nespoulous (CEA Paris-Saclay)

Dr. Lucia Sargentini (CEA Paris-Saclay)

Dr. Riccardo Finotello (CEA Paris-Saclay)

Short-talk

Wednesday, 17 September 2025, 16:00

Critical Heat Flux (CHF) represents a concern for the nuclear safety, as it leads to a rapid drop down in the heat transfer between a heated surface and the liquid coolant in the core of nuclear reactors. This could cause several issues to the system, including structural damage and release of radioactive material.

The main challenge related to CHF prediction is the highly non-linear relationship with the physical features it depends on. For that reason, the prediction of CHF is often affected by large uncertainties.

In this research, the CHF database provided by the U.S. Nuclear Regulatory Commission (NRC) is utilized to develop machine learning (ML) methods for CHF prediction and robust uncertainty quantification (UQ) techniques. The performance of the ML models is assessed against established data-driven strategies, while a coverage-based approach is considered for UQ methods by using conformal prediction and a quality-driven loss function. It is found that it is possible to confidently estimate the prediction uncertainties with a 95% coverage regarding experimental CHF values.

Spotlight Talks

Thursday, 18 September 2025, 9:00 – 11:00

Chair: **Dr. Harold Erbin** (IPhT)

Spotlight talks present cutting-edge research applications of uncertainty quantification in machine learning across various scientific domains. These talks highlight innovative methodologies and their implications for advancing both theoretical understanding and practical applications.

Machine learning of biological sequences

Dr. Jorge Fernandez-de-Cossio-Diaz (IPhT)

Spotlight

Thursday, 18 September 2025, 9:00

Over the last decade machine learning has had tremendous impact on biological sequence data analysis. In this talk, I will begin by introducing general issues related to biological sequence modeling. I will then review a selection of recent works on this topic, including: i) generative models for sequence design, ii) sampling of evolutionary paths between natural sequences of different classes, and iii) predictive models of directed evolution. I will also discuss some sources of uncertainty that arise with biological sequence data in different contexts (alignment, phylogenetic correlations, sampling noise, ...), their potential impact on the models, and efforts to mitigate it.

Interaction reconstruction in scintillator detectors for PET imaging - Deep Learning approach with uncertainty quantification

Dr. Geoffrey Daniel (CEA Paris-Saclay)

Short-talk

Thursday, 18 September 2025, 10:30

Positron Emission Tomography (PET) is a medical imaging modality that is powerful to follow biological processes. Nevertheless, it is of importance to increase its sensitivity to improve the contrast on the PET images and to decrease the patient exposure to radiation. One promising way is the use of the Time-of-Flight (ToF) of coincident gamma ray photons to get a more precise information of the annihilation location event by event. New instrumental developments are necessary to achieve this objective, especially ultra-fast detectors. The ClearMind project has developed a detection system based on a fast lead tungstate (PbWO_4) monolithic scintillator detector and its conception necessitates new processing methods of its signals to reconstruct the gamma photons interactions.

This work focuses on the processing of the recorded signals to estimate the spatial coordinates of the gamma interactions within the detector. The complexity of these signals brings the necessity to use advanced tools and we have developed Deep Learning models, trained on simulation. We introduce a custom loss function that aims at estimating the inherent uncertainties due to the randomness of the signal formation and that incorporates the physical constraints of the detector. The results show the effectiveness of the proposed approach that provides a robust and reliable estimation of the interaction location. They highlight the benefit of the uncertainties estimation that will be exploited in the future for PET image reconstruction to discard or weight each individual event in the objective to improve the signal to noise ratio on the reconstructed image.