T 96: Data, AI, Computing 7 (uncertainties, likelihoods)

Time: Thursday 16:00–18:15 Location: Geb. 30.33: MTI

T 96.1 Thu 16:00 Geb. 30.33: MTI

Effects of adversarial attacks and defenses on generic neural network applications in high energy physics — •Timo Saala and Matthias Schott — Johannes Gutenberg-Universität Mainz

Neural networks have emerged as pivotal tools within high-energy physics (HEP). A field that has recently gained a lot of traction in general deep learning is adversarial learning, which concerns itself with generating adversaries that can be leveraged in order to fool neural networks. Adversaries, intentionally crafted for maximal classification or regression errors with minimal visible input perturbation have spurred the development of techniques for both generating, as well as defending against them. A subset of defense techniques can additionally be applied in order to improve the robustness, and sometimes even the generalization capabilities of deep neural networks. Moreover, adversarial attacks and defenses could potentially offer a means to define the systematic uncertainties in neural networks.

In this study, we employ adversarial learning techniques on multiple neural networks from the HEP environment, reconstructed using exclusively CMS Open Data in order to ensure replicable findings. Through the deployment of adversaries, we not only assess the robustness of these networks but also apply adversarial defenses aiming for the construction of HEP networks displaying larger robustness and better generalization.

T 96.2 Thu 16:15 Geb. 30.33: MTI

Using Adversarial Attacks to Fool IceCube's Deep Neural Networks — •Oliver Janik[1], Philipp Soldin[2], and Christopher Wiebusch[2] — [1]FAU Erlangen-Nürnberg, Germany — [2]RWTH Aachen, Germany

Deep neural networks (DNNs) find more and more use in the data analysis of physics experiments. In the context of adversarial attacks, it has been observed that imperceptible changes to the input of DNNs can alter the output drastically. These adversarial attacks are utilized to investigate DNNs used in particle and astroparticle physics within the AI Safety project. While existing algorithms, like DeepFool, can successfully attack those networks, they produce physically improbable changes. A new method has been developed to vary the inputs only within their uncertainties. The algorithm is applied to an exemplary DNN from the IceCube Neutrino Observatory for particle identification. This network's robustness and unique evaluation prospects are presented using the developed fooling algorithm.

T 96.3 Thu 16:30 Geb. 30.33: MTI

Sharing AI-based Searches with Classifier Surrogates — •Sebastian Bieringer[1], Gregor Kasieczka[1], Jan Kieseler[2], and Mathias Trabs[3] — [1]Universität Hamburg, Institut für Experimentalphysik, Luruper Chaussee 149, 22761 Hamburg, Germany — [2]Karlsruhe Institute of Technology, Institute of Experimental Particle Physics, 76131 Karlsruhe, Germany — [3]Karlsruhe Institute of Technology, Institute of Stochastics, 76131 Karlsruhe, Germany

In recent years, Neural Network-based classification has been used to improve evaluations at collider experiments. While this strategy proofs to be hugely successful, the underlying models are not commonly shared with the public as they are based on experiment-internal simulation. We propose to cater a generative model, a Classifier Surrogate, sampling the classification output from high-level jet information, with every Neural Network-based evaluation to enable further tests on the evaluation. Continuous Normalizing Flows are one suitable generative architecture that can be efficiently trained using Conditional Flow Matching and easily extended by Bayesian uncertainties to indicate the unknown inputs to the user. For a top-tagging example, we demonstrate the application of Flows in combination with uncertainty estimation through either inference of a mean-field Gaussian weight posterior, or Monte Carlo sampling network weights.

T 96.4 Thu 16:45 Geb. 30.33: MTI

Combining data with unknown correlations — •Lukas Koch — JGU Mainz

The combination of data points is a regular necessity in particle physics: be it to calculate an ”average” of multiple measurements of the same thing, or to do model tests and fits to one or more data sets with multiple data points each. Under ideal circumstances, the uncertainties and correlations between all data points – i.e. the joint likelihood function – is known. In that case, it is trivially possible to do ”the right thing” and, e.g., use the Mahalanobis distance – or ”chi-squared” – calculated with the known covariance matrix in order to do statistical tests with the expected properties. In reality, at least some of that information is missing, e.g. when there is no information about the correlation between the results from two separate experiments which share some systematics, or – especially for older publications – when there is no publicly available covariance matrices for an experimental result. Applying the M-distance under the assumption of no correlation can lead to undercoverage in this case. In this talk, I will present the use of alternative test statistics that behave conservative in these circumstances, and thus could be a more robust choice when faced with this issue.

T 96.5 Thu 17:00 Geb. 30.33: MTI

Binary Black Hole Parameter Estimation using a Conditioned Normalizing Flow — •Markus Bachlechner, Oliver Pooth, and Achim Stahl — III. Physikalisches Institut B, RWTH Aachen University

The proposed Einstein Telescope is the first of the third-generation gravitational wave detectors. It is expected to reach a noise level at least an order of magnitude lower than current interferometers like LIGO and Virgo. The thus improved sensitivity increases the observable volume and extends the time window in which the inspiral phase of binary systems is measurable. To analyze the resulting vast amounts of data efficiently, Neural Networks (NNs) can be utilized. This talk presents a fast Binary Black Hole parameter reconstruction by applying a conventional convolutional NN which conditions a subsequent Normalizing Flow (NF). Using the NF, an approximated posterior parameter distribution on an event-by-event basis is obtained, and thus uncertainties can be estimated.

T 96.6 Thu 17:15 Geb. 30.33: MTI

Probabilistic Machine Learning for the XENONnT position reconstruction — •Sebastian Vetter — Karlsruhe Institute of Technology, Institute for Astroparticle Physics

The XENONnT detector is a dual-phase Xenon time projection chamber to search for Dark Matter. To fully exploit background reduction, it is important to know the exact position of events in the detector. The event position reconstruction is commonly performed by a combination of different neural networks (NNs). These NNs, like most machine learning models used in modern experiments, output a singular point in the parameter space. The parameter space in this example is the horizontal plane of the detector.

In this talk I will present and compare two ways of modifying NNs to change their output from a singular point to a probability density.

The resulting probability density functions provide information about the uncertainty of the predictions. The numerical value of the uncertainty can be used to filter for potentially incorrectly reconstructed events. The shape of the uncertainty distribution can be analyzed to learn about trends and biases in the position reconstruction, ultimately leading to an improved signal to background discrimination.

This work is supported in part through the Helmholtz Initiative and Networking Fund (grant agreement no. W2/W3-118). In addition, support by the graduate school KSETA at KIT is gratefully acknowledged.

T 96.7 Thu 17:30 Geb. 30.33: MTI

`dilax`: Differentiable Binned Likelihoods in JAX — •Peter Fackeldey, Benjamin Fischer, Felix Zinn, and Martin Erdmann — III. Physikalisches Institut A, RWTH Aachen University

`dilax` is a software package for statistical inference using likelihood functions of binned data. It fulfils three key concepts: performance, differentiability, and object-oriented statistical model building.

`dilax` is build on JAX - a powerful autodifferentiation Python framework. By making every component in `dilax` a “PyTree”, each component can be jit-compiled (`jax.jit`), vectorized (`jax.vmap`) and differentiated (`jax.grad`). This enables additionally novel computational concepts, such as running thousands of fits simultaneously on a GPU.

We present the key concepts of `dilax`, show its features, and discuss performance benchmarks with toy datasets.

T 96.8 Thu 17:45 Geb. 30.33: MTI

Building and Evaluation of Likelihood Functions with `dilax` – Differentiable Likelihoods in `JAX` — Peter Fackeldey, Benjamin Fischer, •Felix Zinn, and Martin Erdmann — III. Physikalisches Institut A, RWTH Aachen University

A common task in high energy phsics (HEP) is the measurement of physical quantities, such as the cross section of a physics process, using likelihood functions of binned data.

The python software package `dilax` allows to define these likelihood functions. It is purely based on `JAX` and thus enables novel computing concepts such as automatic differentiation and vectorization in the context of likelihood fitting.

In this talk we show how to build and evaluate a likelihood function in `dilax`. We present how to perform a likelihood fit including systematic uncertainties in the context of HEP analyses. The results will be validated with existing fitting libraries commonly used in HEP.

T 96.9 Thu 18:00 Geb. 30.33: MTI

Refining Fast Simulations using Machine Learning Techniques — Samuel Bein[2], Patrick Connor[2], Sebastian Götschel[1], Daniel Ruprecht[1], Peter Schleper[2], •Lars Stietz[1,2], and Moritz Wolf[2] — [1]Technische Universität Hamburg — [2]Universität Hamburg

In the realm of particle physics, a large amount of data are produced in particle collision experiments such as the CERN Large Hadron Collider (LHC) to explore the subatomic structure of matter. Simulations of the particle collisions are needed to analyse the data recorded at the LHC. These simulations rely on Monte Carlo techniques to handle the high dimensionality of the data. Fast simulation methods (FastSim) have been developed to cope with the significant increase of data that will be produced in the coming years, providing simulated data 10 times faster than the conventional simulation methods (FullSim) at the cost of reduced accuracy. The currently achieved accuracy of FastSim prevents it from replacing FullSim. We propose a machine learning approach to refine high level observables reconstructed from FastSim with a regression network inspired from the ResNet approach. We combine the mean squared error (MSE) loss and the maximum mean discrepancy (MMD) loss. The MSE (MMD) compares pairs (ensembles) of data samples. We examine the strengths and weaknesses of each individual loss function and combine them as a Lagrangian optimization problem.