Second MODE Workshop on Differentiable Programming for Experiment Design

Monday 12 September 2022 - Friday 16 September 2022

OAC conference center, Kolymbari, Crete, Greece.

Book of Abstracts

Contents

Opening and Keynote Lectures / 1

Lecture: Differentiable programming and design optimization (TBC)

Corresponding Author: pietro.vischia@cern.ch

MODE Data Challenge / 2

MODE Data Challenge

Corresponding Author: giles.chatham.strong@cern.ch

Opening and Keynote Lectures / 3

Keynote Lecture: JAX/DEX for scientific computing

Opening and Keynote Lectures / 36

Introduction and goals of the Workshop

Corresponding Author: tommaso.dorigo@cern.ch

Wine Tasting and Poster Session / 39

Differentiable Preisach Modeling for Characterization and Optimization of Particle Accelerator Systems with Hysteresis

Authors: Ryan Roussel[None]; Auralee Edelen[None]; Juan Pablo Gonzalez-Aguilera[1]; Kabir Dubey[1]; Young-Kee Kim[2]; Nikita Kuklev[1]; Daniel Ratner[3]

[1] *University of Chicago*

[2] *University of Chicago (US)*

[3] *SLAC*

Corresponding Author: rroussel@slac.stanford.edu

Future improvements in particle accelerator performance are predicated on increasingly accurate online modeling of accelerators. Hysteresis effects in magnetic, mechanical, and material components of accelerators are often neglected in online accelerator models used to inform control algorithms, even though reproducibility errors from systems exhibiting hysteresis are not negligible in high precision accelerators. In this Letter, we combine the classical Preisach model of hysteresis with machine learning techniques to efficiently create nonparametric, high-fidelity models of arbitrary systems exhibiting hysteresis. We experimentally demonstrate how these methods can be used in situ, where a hysteresis model of an accelerator magnet is combined with a Bayesian statistical model of the beam response, allowing characterization of magnetic hysteresis solely from

beam-based measurements. Finally, we explore how using these joint hysteresis-Bayesian statistical models allows us to overcome optimization performance limitations that arise when hysteresis effects are ignored.

Applications in AstroHEP / 40

Towards an End-to-End Optimization of IceCube-Gen2 Neutrino Observatory

Author: Christian Glaser[None]

Corresponding Author: christian.glaser@physics.uu.se

I will present the status of optimizing the radio detector of the planned IceCube-Gen2 neutrino observatory at the South Pole. IceCube-Gen2 will enable neutrino astronomy at ultra-high energies (UHE) and will provide insights into the inner processes of the most violent phenomena in our universe. Detecting these UHE neutrinos would be one of the most important discoveries in astroparticle physics in the 21st century. However, the radio detector has not been optimized for its primary objective of measuring the neutrino's energy and direction because of the limitations of current methods. We plan to solve that using differential programming and deep learning to optimize the detector design end-to-end. The project is timely as the detector design can still be influenced over the next three years before production starts in 2025.
I will present first results from surrogate models to replace the time-consuming MC simulations, as well as DNNs for event reconstruction to quantify the detector resolution.

Progress in Computer Science / 43

Automatic Differentiation in ROOT

Authors: Garima Singh[1]; Jonas Rembser[2]; Lorenzo Moneta[2]; Vassil Vassilev[3]

[1] *Princeton University (US)*

[2] *CERN*

[3] *Princeton*

Corresponding Author: garima.singh@cern.ch

Automatic Differentiation is a powerful technique to evaluate the derivative of a function specified by a computer program. Thanks to the ROOT interpreter, Cling, this technique is available in ROOT for computing gradients and Hessian matrices of multi-dimensional functions.
We will present the current integration of this tool in the ROOT Mathematical libraries for computing gradients of functions that can then be used in numerical algorithms.
For example, we demonstrate the correctness and performance improvements in ROOT's fitting algorithms. We will show also how gradient and Hessian computation via AD is integrated in the main ROOT minimization algorithm Minuit.
We will show also the present plans to integrate the Automatic Differentiation in the RooFit modelling package for
obtaining gradients of the full model that can be used for fitting and other statistical studies.

Progress in Computer Science / 44

Differentiable Programming on relational languages.

Author: Paul Peseux[1]

[1] *Université de Rouen*

Corresponding Author: paul.peseux@gmail.com

Many acquired data are inherently relational: human or automated annotations create high value relational data. For example, a planet might belong to a stellar system which has a star of a certain type. This stellar system belongs to a galaxy which also has its own type (**Fig.1**).

Relational programming languages are the right tools to handle these strong structured data. Our proposition is to also construct and optimize models within those relational languages.

In the talk, I would like to present how Differentiable Programming, when treated as a first class citizen, allows scientists to create and very easily optimize interpretable, debuggable and scalable models in the context of relational languages.

This new fully differentiable optimization pipeline unlocks the use of 'categorical models' for every scientist. Scientists already use applied logistic regression for stellar classification, but with this new tool it is possible to go much further.

> **Indico rendering error**
>
> ```
> Could not include image: [429] Error fetching image
> ```

Fig.1 - Galaxy example database

Applications in Particle Physics / 45

Automatic differentiation for Monte Carlo processes

Authors: Alberto Ramos Martinez[1]; Bryan Zaldivar[2]; Guilherme Telo[3]

[1] *Univ. of Valencia and CSIC (ES)*

[2] *IFT Madrid*

[3] *IFIC*

Corresponding Author: alberto.ramos@cern.ch

Automatic Differentiation (AD) techniques allows to determine the
Taylor expansion of any deterministic function. The generalization of
these techniques to stochastic problems is not trivial. In this work we explore two approaches to
extend the ideas of AD to Monte Carlo processes, one based on reweighting (importance sampling)
and another one based on the ideas from the lattice field theory community (numerical stochastic
perturbation theory using the Hamiltonian formalism). We show
that, when convergence can be guaranteed, the approach based on NSPT is
able to converge to the Taylor expansion with a much smaller variance.

Applications in Particle Physics / 46

Machine learning and optimisation in particle accelerator operation for CERN

Authors: Michael Schenk[1]; Verena Kain[1]

Co-authors: Brennan Goddard [1]; Elias Fernandez-Combarro Alvarez ; Francesco Maria Velotti [1]; Gianluca Valentino [2]; Giovanni Zevi Della Porta [1]; Joel Axel Wulff [3]; Kevin Shing Bruce Li [1]; Leander Grech [2]; Michele Grossi [1]; Mircea-Marian Popa [4]; Nico Madysa [1]; Niky Bruchon [1]; Simon Hirlaender [5]; Sofia Vallecorsa [1]

[1] *CERN*

[2] *University of Malta (MT)*

[3] *Lund University (SE)*

[4] *University Politehnica of Bucharest (RO)*

[5] *IDA Lab, University of Salzburg (A)*

Corresponding Author: michael.schenk@cern.ch

CERN maintains a dense and diverse physics programme with numerous experiments requiring a broad spectrum of beam types to be produced and provided by the accelerator chain. This combined with the requests for higher beam intensities and smaller emittances at an overall better beam quality makes accelerator operation more and more challenging. The way forward is to exploit automation techniques and improved modelling to boost machine flexibility, availability, and beam reproducibility.

Ideally, one would realise optimisation based on closed-form physics models. However, these are not always available or exploitable. One possible workaround is to rely instead on sample-efficient optimisation algorithms. Sample efficiency is key here given the cost of beam time and the impact on the physics experiments during machine setting up.

This contribution introduces various sample-efficient optimisation algorithms and demonstrates their application to accelerator operation at CERN. Among others, numerical optimisation and diverse reinforcement learning algorithms, including model-free and model-based variants as well as energy-based models employing quantum annealing, are discussed together with their application to beam steering tasks, beam optics matching, noise cancellation, or radio-frequency parameter tuning for CERN's accelerator complex.

Another stream of research is modelling hysteresis, eddy currents, complicated magnetic structures, such as the pole-face windings in the CERN Proton Synchrotron and many more with their impact on beam using machine learning techniques, and to build inverse models in a second step to control the accelerators. The first results of this endeavour will also be presented.

Applications in Particle Physics / 47

Challenges in the optimization of the HGCAL Optical Fibre Plant

Authors: André David [1]; Laurent Forthomme [1]; Loukas Gouskos [1]; Pedro Vieira De Castro Ferreira Da Silva [1]; Stefani Coco Mendez [1]

[1] *CERN*

Corresponding Authors: pedro.silva@cern.ch, laurent.forthomme@cern.ch

The Compact Muon Solenoid (CMS) detector at the CERN Large Hadron Collider (LHC) is undergoing an extensive Phase II upgrade program to prepare for the challenging conditions of the High-Luminosity LHC (HL-LHC). As part of this program, a novel endcap calorimeter that uses almost 6M Silicon and Scintillator sensors is foreseen. These sensors will sample the electromagnetic and hadronic particle showers using 47 longitudinal layers with fine lateral granularity and providing 5D measurements of energy, position and timing. An hierarchy of electronics cards and transmission links shall be used to readout and ship the enormous amount of data to the trigger and data acquisition systems. An opto-electrical conversion is used to convert the signals to fast optical fibers with a max. 10.24 Gb/s throughput.

The exact mapping of fibres, connectors and shuffles needed to be used is subject to different constraints: mechanical spacing for splicing, connectors and routing; packing efficiency, minimizing the so-called dark-fibers throughput balancing of the inputs to the FPGAs in the data acquisition system. Some of these constraints impact significantly the performance and the budget of the system. Thus an optimisation of the optical fibre plant has the potential to: 1) reduce the final cost of the system by reducing the number of dark fibres and the need for additional shuffling in the backend electronics and 2) increase the longevity of the detector by reducing the number of breaks, and therefore losses, in the fibre paths.

In this presentation we exemplify this optimization procedure as a use case. Custom-made algorithms have been employed so far without resorting to a formal definition of a loss function quantifying the different constraints. The challenges faced towards constructing such a loss function are also discussed.

Applications in Particle Physics / 48

LHCb ECAL optimization

Author: Alexey Boldyrev[1]

[1] *NRU Higher School of Economics (Moscow, Russia)*

Corresponding Author: alexey.boldyrev@cern.ch

LHCb ECAL optimization is a good use case for the generic problem of comprehensive optimization of the complex physics detector. Pipeline-based approach for LHCb ECAL optimization is established and used to scan parameter space for desired subspace which met needs of the LHCb experiment. Parameters for the calorimeter optimization include technology, granularity, Moliere Radius, timing resolution for calorimetric modules, and modules distribution in the ECAL area. Physics significance of a selected channel vs. cost of the detector can be seen as an performance metric of the detector optimization. Assuming differentiable surrogate implementation, the pipeline-based approach is able to speed up the optimization cycle w.r.t. black-box optimization. The latest improvements in the pipeline-based approach and its ability to determine the necessary degree of detail of a simulation are discussed.

Applications in AstroHEP / 49

Normalizing flows for differentiable expectation values

Author: Thorsten Glüsenkamp[1]

[1] *Universität Erlangen-Nürnberg*

Corresponding Author: thorsten.gluesenkamp@fau.de

I will give a short introduction to (conditional) normalizing flows and discuss why they are essential to calculate "manageable" differentiable expectation values of continuous random variables. I will discuss some examples where this might be useful and end the talk with a github-package that combines some state of the art conditional normalizing flows to be used with minimal manual labor.

Applications in AstroHEP / 50

Dark Matter direct detection experiments

Author: Susana Cebrian[1]

[1] *Universidad de Zaragoza*

Corresponding Author: scebrian@unizar.es

Dark matter particles pervading our galactic halo could be directly detected by measuring their scattering off target nuclei or electrons in a suitable detector. The expected signal from this interaction is rare, demanding ultra-low background conditions, and small energy deposits below tens of keV would be produced, requiring low energy detection thresholds. Many different and complementary techniques are being applied or under consideration, like solid-state cryogenic detectors, Time Projection Chambers based on noble liquids, scintillating crystals and purely ionization detectors using semiconductors or gaseous targets.

The features of this possible dark matter signal and the detection techniques will be briefly summarized and the latest results obtained, mainly constraining the properties of the dark matter candidates under different scenarios, will be discussed together with the prospects for the coming years. Exploring regions of cross sections where neutrinos become an irreducible background could be at reach for large detectors foreseen for the end of this decade and candidates with increasingly lower masses are being investigated thanks to the development of novel technologies to reduce the energy threshold in small detectors. The discrimination of the expected signal from radioactive background or other noise events is a must when designing these experiments and machine-learning techniques are already being applied in some cases.

Applications in Muon Tomography / 52

TomOpt: Differentiable Optimisation of Muon-Tomography Detectors

Authors: Giles Chatham Strong[1]; Tommaso Dorigo[1]

Co-authors: Andrea Giammanco [2]; Federica Fanzago [1]; Max Lamparth ; Maxime Lagrange [3]; Federico Nardi [1]; Haitham Zaraket [4]; Pietro Vischia [2]; Anna Bordignon [5]; Oleg Savchenko ; Jan Kieseler [6]

[1] *Universita e INFN, Padova (IT)*

[2] *Universite Catholique de Louvain (UCL) (BE)*

[3] *CP3 Universite Catholique de Louvain*

[4] *Lebanese University (LB)*

[5] *INFN-Padova*

[6] *CERN*

Corresponding Author: giles.chatham.strong@cern.ch

The recent MODE whitepaper*, proposes an end-to-end differential pipeline for the optimisation of detector designs directly with respect to the end goal of the experiment, rather than intermediate proxy targets. The TomOpt python package is the first concrete endeavour in attempting to realise such a pipeline, and aims to allow the optimisation of detectors for the purpose of muon tomography with respect to both imaging performance and detector budget. This modular and customisable package is capable of simulating detectors which scan unknown volumes by muon radiography, using cosmic ray muons to infer the density of the material. The full simulation and reconstruction chain is made differentiable and an objective function including the goal of the apparatus as well as its cost and other factors can be specified. The derivatives of such a loss function can be back-propagated to each parameter of the detectors, which can be updated via gradient descent until an optimal configuration is reached. Additionally, graph neural networks are shown to be applicable to muon-tomography, both to improve volume inference and to help guide detector optimisation.

*MODE et al. (2022) Toward the End-to-End Optimization of Particle Physics Instruments with Differentiable Programming: a White Paper, arXiv:2203.13818 \[physics.ins-det\]

Progress in Computer Science / 53

Towards Algorithmic Differentiation of GATE/Geant4

Authors: Max Aehle[None]; Nicolas Gauger[1]

[1] *Technische Universität Kaiserslautern*

Corresponding Author: max.aehle@scicomp.uni-kl.de

Simulations of the detection process are an essential step in the assessment of a proposed detector design. Often, complex Monte-Carlo simulators with a wide range of applications are employed for this task, as they provide the most realistic and adaptable computational models for the interactions between particles and the detector. For instance, the Bergen pCT collaboration develops a digital tracking calorimeter for proton computed tomography (pCT) using the program GATE based on Geant4.

In the context of gradient-based optimization and differentiable programming, the size, complexity and heterogeneity of such codes makes it difficult to integrate common algorithmic differentiation (AD) tools based on operator overloading. As an example, parts of Geant4 v11.0.0 reinterpret G4double's as integers or cast them to (C++) double's, and do therefore not compile when G4double is redefined as an AD type. We believe that such issues are not fundamental and can be overcome with enough work. Yet, they provide a strong motivation to build AD tools that, from the perspective of a user, need less integration efforts.

In this talk, we present such an alternative approach, and the current state of developing it further to support GATE/Geant4.

Progress in Computer Science / 54

Quantification of Uncertainties in a Model-Based Iterative Reconstruction Algorithm for Proton Computed Tomography

Author: Lisa Kusch[1]

Co-authors: Max Aehle ; Nicolas Gauger [2]

[1] *TU Kaiserslautern*

[2] *Technische Universität Kaiserslautern*

Corresponding Author: lkusch@rhrk.uni-kl.de

Proton computed tomography (pCT) is a medical imaging modality with the potential to improve the accuracy of treatment planning for proton-beam radiotherapy. It produces a three-dimensional image of the relative stopping power (RSP) distribution inside an object, given a list of positions and directions of protons before and after passing through the object, along with the corresponding energy losses. The underlying tomographic reconstruction task can be understood as a least-squares problem, and solved by a model-based iterative process.
Various uncertainties due to calibration errors, measurement errors or errors in the track reconstruction process may arise as inputs to the RSP reconstruction. It is crucial to consider the influence of these uncertainties in the modeling of the reconstruction and in the process of designing the tracking calorimeter as a long-term scope of this work. For this purpose, our aim is to investigate on the effect of local perturbations by means of a sensitivity analysis as well as on the effect of random inputs with the help of strategies for uncertainty propagation in a probabilistic description.

In this work, we will present first investigations for both types of uncertainties. Regarding local perturbations, one may observe that the algorithmic description of the reconstruction process is only piecewise differentiable. Jumps arise from the discrete computation of the set of voxels intersected by a proton path. To represent uncertainties in the proton positions in front of the calorimeter, we model the positions as normally distributed random variables, and employ sampling strategies for propagation. Uncertainties arising from the calibration of the magnets may be modelled with a moderate number of random variables. Here, we make use of discrete projection methods to avoid the computationally high costs of sampling. Uncertainties in the measurements may be specifically adressed in modeling when using an extended most likely path formalism. The modeling and propagation of uncertainties allows us to assess the overall impact of this approach on the quality of the reconstruction.

Progress in Computer Science / 55

In-Situ Range Verification for Particle Therapy Based on Secondary Charged Particles

Author: Alexander Schilling[1]

Co-authors: Nicolas Gauger [2]; Ralf Keidel [3]

[1] *University of Applied Sciences Worms*

[2] *Technische Universität Kaiserslautern*

[3] *Fachhochschule Worms (DE)*

Corresponding Author: alexander.georg.schilling@cern.ch

Radiation therapy using protons or heavier ions is sensitive to range errors caused by misalignment of the patient, changes in patient anatomy, and uncertainties in treatment planning. It is therefore of the utmost importance to ensure treatment quality through range verification. Determining the position of the Bragg peak inside the patient can be done through various means such as prompt gamma detection. When using heavier ions, charged secondary particles can also be used for range verification. Protons on the other hand do not produce charged secondaries of high enough energy to be detectable outside the patient. However, the unique properties of the Bergen pCT detector enable proton range verification through secondary charged particles without the need of an additional device in the treatment room. Preliminary results on Monte Carlo data for water phantoms and a pediatric head phantom suggest feasibility for this approach. The detected tertiary charged particles in the digital tracking calorimeter form a point cloud, from which features can be extracted and fed into machine learning models for predicting the range of the initial proton beam.

Progress in Computer Science / 56

Towards Neural Charged Particle Tracking in Proton Computed Tomography with Reinforcement Learning

Author: Tobias Kortus[None]

Co-authors: Nicolas Gauger [1]; Ralf Keidel [2]

[1] *Technische Universität Kaiserslautern*

[2] *Fachhochschule Worms (DE)*

Corresponding Author: kortus@ztt.hs-worms.de

Proton computed tomography (pCT) is an emerging imaging modality, receiving increasing importance for treatment planning in hadron therapy. One of the core challenges in the reconstruction

pipeline for the high-granularity Bergen pCT Digital Tracking Calorimeter (DTC) prototype is the efficient reconstruction of high-multiplicity proton traces throughout the scanner which, in contrast to X-rays used in conventional CT, do not travel in a straight line. For this application, we propose a novel sequential tracking scheme based on model-free reinforcement learning working on graph-structured data, allowing for ground-truth free training of deep neural network architectures. For optimization, we consider only partial-information in terms of a scalar reward function, modeled based on the Gaussian approximation of the underlying theory of multiple Coulomb scattering. Given this information, we aim to find a policy that maximizes the plausibility of reconstructed tracks. We demonstrate on Monte Carlo simulated data the high potential as well as the current limitations of this approach.

Applications in AstroHEP / 57

Optimization of the Future P-ONE Neutrino Telescope

Author: Christian Haack[1]

[1] *Technical University of Munich*

Corresponding Author: christian.haack@tum.de

P-ONE is a planned cubic-kilometer-scale neutrino detector in the Pacific ocean. Similar to the successful IceCube Neutrino Observatory, P-ONE will measure high-energy astrophysical neutrinos to help characterize the nature of astrophysical accelerators. Using existing deep-sea infrastructure provided by Ocean Networks Canada (ONC), P-ONE will instrument the ocean with optical modules - which host PMTs as well as readout electronics - deployed on several vertical cables of about 1km length. While the hardware design of a first prototype cable is currently being finalized, the detector geometry of the final instrument (up to 70 cables) is not yet fixed.
In this talk, I will present the progress of optimizing the detector design using ML-based surrogate models, which replace computationally expensive MC simulations, and, by providing gradients, allow efficient computation of the Fisher Information Matrix as an optimization target.

Applications in Particle Physics / 58

Phase Space Reconstruction of Beams Using Machine Learning Based Representations

Authors: Auralee Edelen[None]; Christopher Mayes[None]; Daniel Ratner[1]; Juan Pablo Gonzalez-Aguilera[2]; Ryan Roussel[None]

[1] *SLAC*

[2] *University of Chicago*

Corresponding Author: rroussel@slac.stanford.edu

Characterizing the phase space distribution of particle beams in accelerators is often a central part of accelerator operation and optimization. However, with the introduction of more exotic target distributions that are ideal for current accelerator applications conventional diagnostic techniques, which assume simplified beam distributions, start to become inaccurate. By contrast, methods that enable high-fidelity phase space distribution reconstructions require specialized procedures or diagnostics that may not be efficient or available for most accelerator facilities. In this Letter, we introduce a machine learning based technique to reconstruct 6D phase space distributions from common accelerator diagnostics. We first demonstrate that our algorithm reconstructs exotic phase space distributions with corresponding confidence intervals in simulation. We then demonstrate the technique in experimental conditions, improving upon results obtained from conventional diagnostic methods. Finally, we discuss how this technique is widely applicable to almost any accelerator diagnostic, enabling simplified 6D phase space reconstruction diagnostics in the future.

Applications in Particle Physics / 59

Towards a differentiable sampling ECAL model with optimal absorber material distribution

Authors: Shivay Vadhera[None]; Fabricio Andres Jimenez Morales[None]; Vincent Boudry[None]

Corresponding Author: shivay.vadhera@cern.ch

The distribution of material in sampling electromagnetic calorimeters has a direct impact on the energy resolution that they can achieve. R&D calorimetry efforts that are inspired for future collider detector concepts aim for the best energy performance given constraints, e.g. mechanical and cost. We address the energy resolution optimization in the case of a Silicon-Tungsten Electromagnetic Calorimeter (SiW-ECAL) prototype where a particle impacts at different energies, a case study for which we have prepared Geant4 simulations. The first step we are taking towards optimizing the Tungsten (absorber material) distribution for a fixed amount and number of layers is to implement a generative (differentiable) model that can replace the available simulated datasets in the optimization pipeline and is able to interpolate in the input parameter space; this will allow for an automatically differentiable ECAL model. We show and discuss our on-going work in the generative models explored for such purpose and the prospects for including them in a more general optimization scenario (i.e. adding further constraints and figures of merit, besides the energy resolution), as well as the differences with respect to more traditional optimization approaches.

Applications in AstroHEP / 60

Muon identification and gamma/hadron discrimination using compact single-layered water Cherenkov detectors powered by Machine Learning techniques

Author: B.S. González[1]

Co-authors: R. Conceição [1]; A. Guillén [2]; M. Pimenta [1]; B. Tomé [1]

[1] *LIP/IST*

[2] *UGR*

Corresponding Author: borjasg@lip.pt

The muon tagging is an essential tool to distinguish between gamma and hadron-induced showers in wide field-of-view gamma-ray observatories. In this work, it is shown that an efficient muon tagging (and counting) can be achieved using a water Cherenkov detector with a reduced water volume and multiple PMTs, provided that the PMT signal spatial and time patterns are interpreted by an analysis based on Machine Learning (ML).

Following the same rationale, the developed analysis has been tested for two different WCD configurations with 4 and 3 PMTs. The output of the ML analysis, the probability of having a muon in the WCD station, has been used to notably discriminate between gamma and hadron induced showers with $S/\sqrt{B} \sim 4$ for shower with energies $E_0 \sim 1\,\text{TeV}$. Finally, for proton-induced showers, an estimator of the number of muons was built by means of the sum of the probabilities of having a muon in the stations. Resolutions about 20% and a negligible bias are obtained for vertical showers with $N_\mu > 10$.

Applications in Muon Tomography / 61

Detection of material contrast in muon tomography data with machine learning

Author: Konstantin Borozdin[1]

[1] *Sandia Research Center*

Corresponding Author: sandia.rcenter@gmail.com

Muon tomography applications often require detection of material contrast. One example of such application is the detection of contraband at the border. Another example is the detection of steel rebars inside reinforced concrete blocks. Sensitivity of material discrimination depends on the detector configuration, exposure time and clutter. We explore how machine learning techniques can be used to improve the muon imaging in general, and material discrimination in particular. We apply a variety of machine learning methods, including different types of regressions and image classification. Potential of these methods and their relative efficiency is discussed.

Progress in Computer Science / 62

Synthesization of Fast Gradients with Enzyme

Authors: Ludger Paehler[1]; Tim Gymnich[1]; William Moses[2]; Valentin Churavy[2]

[1] *Technical University of Munich*

[2] *Massachusetts Institute of Technology*

Corresponding Author: wmoses@mit.edu

The synthesization of fast reverse- and forward-mode gradients is the key to many algorithms in scientific computing such as meta-learning, optimization, uncertainty quantification, stability analysis, machine learning-accelerated simulations, and end-to-end learning. Enzyme is an automatic differentiation tool which in difference to other tools is built into the LLVM compiler toolchain to be able to benefit from optimizations of the compiler synthesizing gradients for the optimized compiler representation of programs to unlock the significant performance gains this affords. Thus generating high performance gradients of any language going to the LLVM intermediate representation such as C/C++, Fortran, Julia, Rust, and others.
In this talk we will present how Enzyme is able to generate these gradients for forward- and reverse-mode for scientific programs which contain GPU-code (CUDA & ROCm) and parallelism (Julia tasks, MPI & OpenMP). In addition we will present how Enzyme is able to vectorize computation inside and outside of differentiation to further accelerate scientific workflows.

Progress in Computer Science / 63

Speeding up differentiable programming with a Computer Algebra System

Author: Remco de Boer[1]

[1] *Ruhr University Bochum*

Corresponding Author: remco.deboer@rub.de

In the ideal world, we describe our models with recognizable mathematical expressions and directly fit those models to large data sample with high performance. It turns out that this can be done with a CAS, using its symbolic expression trees as template to computational back-ends like JAX. The CAS can in fact further simplify the expression tree, which results in speed-ups in the numerical back-end.
The ComPWA project offers Python libraries that use this principle to formulate large expressions

for amplitude analysis, so that the user has the flexibility to quickly implement different formalisms and can also easily perform fast computations on large data samples.

Applications in Particle Physics / 64

Optimization of small scale experiments, MILLIQAN and SUB-MET as examples

Author: Haitham Zaraket[1]

[1] *Lebanese University (LB)*

Corresponding Author: haitham.zaraket@cern.ch

Several new experiments that looks for millicharged particle has been proposed recently, Milliqan is an international collaboration that is now installed and under upgrade at point 5 @ CMS at CERN for a second run. SUBMET is a new proposal that is expected to be installed at JPARC. The common denominator of the two proposals is the simplicity of the construction. Muons are the main background that should be well characterized, this is done through additional side and top panels. Looking for detector optimization constitutes a simple benchmark for detector optimization tools like Tomopt.

Applications in Muon Tomography / 65

Ultra-fast muon tomography simulations using Adversarial Neural Networks

Author: Pablo Martinez Ruiz Del Arbol[1]

[1] *Universidad de Cantabria and CSIC (ES)*

Corresponding Author: pablo.martinez@cern.ch

This presentation explores the possibility of using Generative Adversarial Neural Networks (GANN) in order to simulate the propagation of muons through material without using a complete simulation of the physical processes. In order to achieve this goal, Generative Adversarial Neural Networks have been used to simulate muon tomography data applied to the measurement of the thickne of isolated pipes. The GANNs have shown an excellent ability to reproduce the scattering distributions, even when the thickness of the pipe is provided as a conditional parameter to the network. Studies on the interpolation ability of the network have been also conducted, being very successfull.

Wine Tasting and Poster Session / 66

(Poster) Muon identification and gamma/hadron discrimination using compact single-layered water Cherenkov detectors powered by Machine Learning techniques

Author: BS González[1]

Co-authors: R. Conceição [1]; A. Guillén [2]; M. Pimenta [1]; B. Tomé [1]

[1] *LIP/IST*

[2] *UGR*

Corresponding Author: borjasg@lip.pt

The muon tagging is an essential tool to distinguish between gamma and hadron-induced showers in wide field-of-view gamma-ray observatories. In this work, it is shown that an efficient muon tagging (and counting) can be achieved using a water Cherenkov detector with a reduced water volume and multiple PMTs, provided that the PMT signal spatial and time patterns are interpreted by an analysis based on Machine Learning (ML).

Following the same rationale, the developed analysis has been tested for two different WCD configurations with 4 and 3 PMTs. The output of the ML analysis, the probability of having a muon in the WCD station, has been used to notably discriminate between gamma and hadron induced showers with $S/\sqrt{B} \sim 4$ for shower with energies $E_0 \sim 1\,\mathrm{TeV}$. Finally, for proton-induced showers, an estimator of the number of muons was built by means of the sum of the probabilities of having a muon in the stations. Resolutions about 20% and a negligible bias are obtained for vertical showers with $N_\mu > 10$.

Wine Tasting and Poster Session / 67

Decrypting hadronic showers for studies of particle ID with Deep Learning

Authors: Rumman Neshat[1]; Tommaso Dorigo[2]

[1] *Indian Institute Of Science Education and Research(IISER) kolkata*

[2] *Universita e INFN, Padova (IT)*

Corresponding Authors: tommaso.dorigo@cern.ch, rn18ms035@iiserkol.ac.in

As the first step in a wide-ranging study to determine the capabilities of fine-grained calorimeters to identify different hadrons within dense showers, we show how to extract all the information about all intermediate processes taking place within the development of complex hadron showers produced by simulation in GEANT4.

Applications in Muon Tomography / 68

Deep neural networks applied to muon tomography scattering angular distributions

Author: Carlos Diez[1]

[1] *Muon Tomography Systems S.L.*

Corresponding Author: carlos@muon.systems

Several applications of scattering muon tomography require the estimation of a limited number of key parameters associated to a given sample. In this presentation we explore the use of the quantiles of the angular and spatial deviation distributions as the input to Deep Neural Networks regressing on the parameters of interest. We provide examples related to the measurement of the position of the electrodes in an electric furnace and also on the estimation of the snow water equivalent (SWE) in snowpacks.

Wine Tasting and Poster Session / 69

Muon momentum measurement in Muon Scattering Tomography

Author: Maxime Lagrange[1]

Co-authors: Andrea Giammanco [2]; Anna Bordignon ; Federica Fanzago [3]; Federica Fanzago [4]; Giles Chatham Strong [4]; Haitham Zaraket [5]; Jan Kieseler [6]; Max Lamparth ; Oleg Savchenko ; Pietro Vischia [2]; Tommaso Dorigo [4]

[1] *CP3 Universite Catholique de Louvain*

[2] *Universite Catholique de Louvain (UCL) (BE)*

[3] *Unknown*

[4] *Universita e INFN, Padova (IT)*

[5] *Lebanese University (LB)*

[6] *CERN*

Corresponding Author: maximelagrange98@gmail.com

The recent MODE whitepaper*, proposes an end-to-end differential pipeline for the optimization of detector designs directly with respect to the end goal of the experiment, rather than intermediate proxy targets. The TomOpt python package is the first concrete step in attempting to realize such a pipeline, and aims to allow the optimisation of detectors for the purpose of muon tomography with respect to both imaging performance and detector budget. Within this context, we propose to add an extra detection module in order to measure muon momentum. Given that momentum knowledge improves the quality of image reconstruction, a trade-off has to made between the cost of the momentum measurement module and the imaging performance required by the MST task. This poster explores how parameters of the detector can be related to the overall performance of the MST detection system.

Applications in AstroHEP / 70

The Einstein Telescope project

Author: Giacomo Bruno[1]

[1] *Universite Catholique de Louvain (UCL) (BE)*

Corresponding Author: giacomo.bruno@cern.ch

The Einstein Telescope (ET) is the future European terrestrial gravitational wave observatory based on laser interferometry. The project is currently in the preparatory phase after its integration in the EU ESFRI roadmap for large research infrastructures in 2021 and the official creation of the scientific Collaboration in 2022. A number of key design innovations are foreseen for ET like underground operation, triangular shape, longer arms and use of a pair of complementary interferometers each optimised in different frequency range. An intense research and development program on ET technology is ongoing. This talk will give an overview of the ET science as well as its design and technology challenges.

Opening and Keynote Lectures / 72

Keynote Lecture: High-performance Algorithmic Differentiation

Corresponding Author: max.sagebaum@scicomp.uni-kl.de

Applying algorithmic differentiation (AD) to computer code is a challenging task. To this end, several tools have been developed over the last decade, which ease the application of AD. With MeDiPack, it is no longer a problem to handle all kinds of MPI communication. The recently released library OpDiLib provides out of the box AD capabilities for OpenMP parallel codes. Modern AD tools, like CoDiPack, make use of expression templates to improve the taping speed and memory footprint of the differentiated code. This makes the application of AD to computer code much more applicable. The talk will show use cases where AD has been applied to large scale applications and on these uses cases the techniques for applying AD will be discussed. Therefore, we will look into the tools MeDiPack, OpDiLib as well as CoDiPack and why they have been implemented and how they solve the usual problems of applying AD. For AD tools, a general overview is given on the implementation strategies and their advantages and drawbacks. With this information a more elaborated comparison on AD tools can be performed and criteria for an efficient AD tool can be determined. From these criteria CoDiPack, has been developed. One question that is seldom addressed in applying AD to a code, is how the code can be validated. We will look at techniques how this can be done and develop from these insights a framework that can be used to not only validate the differentiated code but also to analyse its AD performance. This can then be used to further improve the differentiated code. Finally, a conclusion is drawn and we will elaborate what the next generation of AD tools needs to provide for more efficiency.

Wine Tasting and Poster Session / 73

Concept of measuring (Multi-)Strange Hadron Yields in the CBM Experiment using Machine Learning Techniques

Author: Axel Puntke[1]

[1] *WWU Münster*

Corresponding Author: axel.puntke@uni-muenster.de

The Compressed Baryonic Matter (CBM) experiment at FAIR will investigate the QCD phase diagram at high net-baryon density ($\mu B > 500$ MeV) with heavy-ion collisions in the energy range of $\sqrt{s}NN = 2.7$–4.9 GeV. Precise determination of dense baryonic matter properties requires multi-differential measurements of strange hadron yields, both for the most copiously produced K0s and Λ as well as for rare (multi-)strange hyperons and their antiparticles.

The strange hadrons are reconstructed via their weak decay topology using PFSimple, a Kalman Filter Mathematics-based package that has been developed for the reconstruction of particles via their weak decay topology. The large combinatorial background needs to be removed by applying certain selection criteria to the topological features.

In this poster, selection criteria optimization for strange hadrons using the boosted decision tree-based library XGBoost will be discussed and the performance of this non-linear multi-parameter selection method is evaluated. To gain insights into the importance of the different features, the trained model is analyzed by looking at the SHAP values, which give an overview of the impact of a given feature value on the prediction and can help to improve the accuracy of the model. As the CBM experiment is under construction and therefore no real data is available yet, signal and background data for the machine learning model are both taken from simulated data of two different event generators: DCM-QGSM-SMM generates the signal candidates and UrQMD data is treated as real data.

Wine Tasting and Poster Session / 74

Reconstruction techniques for nuclear reactions measured with ACtive TARget TPC

Author: Lorenzo Domenichetti[1]

Co-authors: Thomas Roger [2]; Tommaso Marchi [3]

[1] *INFN LNL, University of Padua*

[2] *GANIL*

[3] *INFN LNL*

Corresponding Author: lorenzo.domenichetti@lnl.infn.it

ACTAR is an active-target TPC optimized for the study of nuclear reactions produced by low-intensity beams, such as radioactive beams. In this detector, the gas used to track charged particles within the chamber is at the same time used as a target for the incoming beam. Reconstructing the tracks left by ions is a challenging task, and two different reconstruction algorithms are compared in this work. The first is based on the RANSAC paradigm, a random and iterative routine, while the second is related to the Hough transform. These different approaches will be exposed and their performance compared, based on experimental and simulated data. In addition, results from the application of Convolutional Neural Networks to ACTAR data will be presented, focusing on both the advantages brought by this technique and the challenges still to be faced.

Wine Tasting and Poster Session / 75

Automatic Optimization of Derivatives with Polyhedral Compiler

Author: Alexander Demin[1]

Co-authors: Chris Elrod [2]; Yingbo Ma [2]

[1] *Max Planck Institute for Informatics*

[2] *Julia Computing*

Corresponding Author: alexander.demin@mpi-inf.mpg.de

Derivatives, mainly in the form of gradients and Hessians, are ubiquitous in machine learning and Bayesian inference. Automatic differentiation (AD) techniques transform a program into a derivative (adjoint) program, which is run to compute the gradient.

Traditionally, most AD systems have been high level, and unable to extract good performance on scalar code or loops modifying memory. These have instead been reliant on extensive libraries of optimized kernels –e.g. for linear algebra, convolutions, or probability distributions –combined with associated adjoint rules, forcing practitioners to express their models in terms of these kernels to attain satisfactory performance.

New low level AD tools such as Enzyme allow for differentiating kernels implemented as naive loops. However, since the derivative code is generated programmatically during the application of reverse mode AD, the reverse passes are likely to access memory in patterns suboptimal with respect to both cache and SIMD performance. This opens a new opportunity for polyhedral loop-and-kernel compilers to provide the aggressive transforms needed for high performance.

We report on work in progress on the implementation of LoopModels –an automatic loop optimization library based on LLVM IR using polyhedral dependency analysis methods to attain performance competitive with vendor libraries on many challenging loop nests, with applications from linear algebra to the derivative code produced by AD.

Wine Tasting and Poster Session / 76

A flexible and efficient machine learning approach for data quality monitoring

Authors: Andrea Wulzer[1]; Gaia Grosso[2]; Jacopo Pazzini[3]; Marco Letizia[None]; Marco Rando[4]; Marco Zanetti[2]; Nicolò Lai[5]

[1] *CERN and EPFL*

[2] *Universita e INFN, Padova (IT)*

[3] *Università e INFN, Padova (IT)*

[4] *Universita' degli Studi di Genova*

[5] *Università di Padova*

Corresponding Author: marco.letizia@ge.infn.it

We present a machine learning approach for real-time detector monitoring. The corresponding core algorithm is powered by recent large-scale implementations of kernel methods, nonparametric learning algorithms that can approximate any continuous function given enough data. The model evaluates the compatibility between incoming batches of experimental data and a reference data sample, by implementing a hypothesis testing procedure based on the likelihood ratio. The resulting model is fast, efficient and agnostic about the type of potential anomaly in the data. We show the performance of the model on multivariate data from muon chamber monitoring.

Applications in AstroHEP / 77

Experiment design for 21st Century astronomy

Author: Emille Eugenia DE OLIVEIRA ISHIDA[1]

[1] *CNRS*

Corresponding Author: emille.ishida@clermont.in2p3.fr

The recent increase in volume and complexity of astronomical data has fomented the use of machine learning techniques. However, the acquisition of labels in astronomy is, by construction, very expensive and time-consuming. In this context, experiment design tasks are aimed at optimizing the allocation of scarce labelling resources. The proper application of such methods will be crucial to ensure optimum scientific potential exploitation from the upcoming Vera Rubin Observatory Large Survey of Space and Time (LSST) – expected to start operations in 2024. This task needs to take into account not only the number of hours per night available in telescopes dedicated to the labeling task (knapsack constraints), but also a series of observational conditions which change the cost of each label from one night to the next. In this talk I will describe the astronomical data scenario and the challenges it poses for traditional experiment design strategies with the context of LSST. As a case study, I will describe efforts of developing a full adaptive learning pipeline for telescope time allocation, show how it impacts scientific results and discuss the challenges ahead.

Applications in AstroHEP / 78

Challenges in Cosmic Microwave Background Data Analysis

Corresponding Author: barreiro@ifca.unican.es

The Cosmic Microwave Background (CMB) is a radiation that reaches us from all the directions of the sky. It originated shortly after the Big Bang and is the oldest radiation that we can observe in the

Universe, providing us with very valuable information about the early universe and how it evolved. CMB radiation is also polarized. In particular, the next big challenge in the CMB field is to detect the so-called primordial B-mode of polarization, a very weak signal sourced by a Primordial Background of Gravitational Waves, as predicted by inflation. Therefore, its detection would constitute a very solid proof of this theory as well as a major discovery in Physics.

However, CMB observations also contain a number of contaminant astrophysical signals and instrumental noise, which are mixed in the data and that must be separated from the CMB before deriving any meaningful cosmological conclusion. This is particular challenging since the signal of interest is much weaker than the rest of the components and the characteristics of the contaminant signals are not always well known.

Therefore, a key point of CMB data analysis is the so-called component separation process that aims to extract the cleanest possible CMB signal. I will present the most common approaches to carry out component separation and discuss the possibility of tackling this problem using machine learning, whose application has recently started in this field.

Applications in Muon Tomography / 79

Detector optimization in Muon scattering Tomography

Author: Maxime Lagrange[1]

Co-authors: Andrea Giammanco [2]; Anna Bordignon ; Federica Fanzago [3]; Federica Fanzago [4]; Giles Chatham Strong [4]; Haitham Zaraket [5]; Jan Kieseler [6]; Max Lamparth ; Oleg Savchenko ; Pietro Vischia [2]; Tommaso Dorigo [4]

[1] *CP3 Universite Catholique de Louvain*

[2] *Universite Catholique de Louvain (UCL) (BE)*

[3] *Unknown*

[4] *Universita e INFN, Padova (IT)*

[5] *Lebanese University (LB)*

[6] *CERN*

Corresponding Author: maximelagrange98@gmail.com

Muon Scattering Tomography (MST) has been through a fruitful period of development which led to many applications in various fields such as border controls, nuclear waste characterization, nuclear reactor monitoring and preventive maintenance of industrial facilities. Whatever the use case, MST detector conception aims at reaching the best performance that respects specific constraints, whether in terms of budget, or data acquisition time. In this context, the recent work of the MODE collaboration led to the creation of TomOpt, an end-to-end differentiable pipeline aiming at optimizing muon scattering tomography detectors conception.

This presentation will illustrate which parameters of an MST experiment, from the detector geometry and technology to hidden parameters of image reconstruction and material classifier algorithms, can be included in such an optimization. I will give hints on how to evaluate the cost and performance of a MST tracking system based on its conception features. The inclusion of muon energy spectrometer within border guard MST experiments and its effect on imaging performance will also be discussed.

Applications in Particle Physics / 81

Score-function Optimization⊠ for Branching Processes

Corresponding Author: lukas.heinrich@cern.ch

One of the key difficulties in making HEP differentiable is the highly stochastic and discrete nature of both simulation and reconstruction. While not directly differentiable, gradients of expectation values of stochastic simulator output can be estimated using probabilistic programming and score functions. In this talk I will demonstrate score function based optimization of material maps on a toy particle shower example and discuss possible future directions.

Applications in Nuclear Physics / 82

Adaptive Experimentation to assist the detector design of the future Electron Ion Collider

Corresponding Author: cfanelli@wm.edu

The Electron Ion Collider (EIC), the future ultimate machine to study the strong force, is a large-scale experiment with an integrated detector that covers the central, far-forward, and far-backward regions. EIC is utilizing AI starting from the design phase in order to deal with compute intensive simulations and a design made by multiple sub-detectors — each characterized by multiple design parameters and driven by multiple design criteria.
In this context, AI offers state of the art solutions to design the experiment in a more efficient way. This talk provides an overview of the recent progress made during the EIC detector proposal; it will also cover how this work could further progress in the near future.

83

Music Offer (piano concert)

Corresponding Author: tommaso.dorigo@cern.ch

Organization of Future Activities and Closing of the Workshop / 84

Summary of the workshop and closing words

Corresponding Author: pietro.vischia@cern.ch

Applications in Particle Physics / 85

Surrogate Regressors for a Surrogate generators

Authors: Alexander Rogachev[1]; Fedor Ratnikov[2]

[1] *National Research University Higher School of Economics (RU), Yandex School of Data Analysis (RU)*

[2] *Yandex School of Data Analysis (RU)*

Corresponding Author: fedor.ratnikov@gmail.com

High energy physics experiments essentially rely on the simulation data used for physics analyses. However, running detailed simulation models requires a tremendous amount of computation resources. New approaches to speed up detector simulation are therefore needed.

t has been shown that deep learning techniques especially Generative Adversarial Networks may be used to reproduce detector responses. However, those applications are challenging, as the generated responses need evaluation not only in terms of image consistency, but different physics-based quality metrics should be taken into consideration as well.

In our work, we develop a multitask GAN-based framework with the goal to speed up the response generation of the Electromagnetic Calorimeter. We introduce the Auxiliary Regressor as a second task to evaluate a proxy metric of the given input that is used by the Discriminator of the GAN. We show that this approach improves the stability of GAN and the model produces samples with the physics metrics distributions closer to the original ones than the same architecture without additional loss.

Applications in Muon Tomography / 86

NEWCUT: a novel tool for reconstruction of low-energy muon spectra

Corresponding Author: olah@g.ecc.u-tokyo.ac.jp

High-precision muographic imaging of targets with thicknesses up to a few meters requires the accurate measurement of the angular and energy dependent flux of cosmic-ray muons in the low-energy regime (up to a few GeV). We designed a muon spectrometer, called NEWCUT 1. It is a six-meter-length tracking system consists of nineteen Multi-wire Proportional Chambers (MWPCs) and lead plates. The rotatable support structure allows to tracking the charged particle between the horizontal and vertical directions. A machine learning-assisted muon energy classifier was implemented and trained using coordinate and energy deposit data simulated on chamber-by-chamber in GEANT4 framework. The comparison of simulated and reconstructed muon spectra suggest that the actual arrangement of NEWCUT is applicable to measure the muon spectra up to an energy of 6 GeV. The data analysis methods and first experimental results will be discussed.

1 Oláh et al. Development of Machine Learning-Assisted Spectra Analyzer for the NEWCUT Muon Spectrometer, Journal of Advanced Instrumentation in Science, vol. 2022, https://doi.org/10.31526/jais.2022.264 (2022).

Progress in Computer Science / 87

Differentiable Programming and the IRIS-HEP Analysis Grand Challenge

Corresponding Author: gwatts@uw.edu

The IRIS-HEP Analysis Grand Challenge (AGC) seeks to build and test a fully representative HL-LHC analysis based around new analysis tools being developed within IRIS-HEP and by others. The size of the HL-LHC datasets is expected to require fully distributed analyses sometimes sourcing 100's of TB of data or event, later in the HL-LHC, PB-sized datasets. This talk will give a brief overview of the AGC and discuss how differentiable programming will be incorporated into the challenge as a stretch goal.

Opening and Keynote Lectures / 88

Presentation by the Head of the Institute of Theology and Ecology - Department of the Orthodox Academy of Crete

Opening and Keynote Lectures / 89

Welcome by the General Director of the Orthodox Academy of Crete

Applications in Muon Tomography / 90

Scattering Tomography with the EM Algorithm

Corresponding Author: anna.bordignon.2@studenti.unipd.it

We investigate the problem of distinguishing materials by scattering tomography with an R implementation of the EM (Expectation Maximization) algorithm and we compare the results with those from the PoCA (Point of Closest Approach) method.

Organization of Future Activities and Closing of the Workshop / 91

Prize awards for poster competition

Author: The Poster Committee[None]

Corresponding Authors: roberto.ruiz.de.austri@cern.ch, gwatts@uw.edu, lkusch@rhrk.uni-kl.de

MODE Data Challenge / 92

Prize awards for data challenge competition

Applications in Nuclear Physics / 93

FAST EMULATION OF DEPOSITED DOSE DISTRIBUTIONS BY MEANS OF DEEP LEARNING

Authors: Lorenzo Arsini[1]; Andrea Ciardiello[2]; Barbara Caccia[3]; Angelica De Gregorio[4]; Riccardo Faccini[5]; Gaia Franciosini[4]; Stefano Giagu[6]; Annalisa Muscato[7]; Carlo Mancini Terracciano[1]

[1] *Department of Physics, University of Rome "La Sapienza", Rome, Italy, NFN, Section of Rome, Rome, Italy*

[2] *INFN, Section of Rome, Rome, Italy*

[3] *Istituto Superiore di Sanità, Rome, Italy*

[4] *INFN - National Institute for Nuclear Physics*

[5] *Department of Physics, University of Rome "La Sapienza", Rome, Italy*

[6] *Sapienza Universita e INFN, Roma I (IT)*

[7] *Scuola post-laurea in Fisica Medica, University of Rome "La Sapienza", Rome, Italy*

Corresponding Author: lorenzo.arsini@uniroma1.it

Background: The aim of our work is to build a Deep Learning algorithm capable of generating the distribution of absorbed dose by a medium interacting with a given particle beam.

Such an algorithm can provide a precise and faster alternative to a Monte Carlo (MC) simulation, which are currently employed in the optimisation process of radio therapy treatment (RT) planning. A faster dose data generation can be beneficial for every kind of RT treatment planning and especially relevant in contexts, as electron FLASH RT 1 or Volumetric Modulated Arc Therapy (VMAT), that offer much more freedom in the choice of the entry angles of the beam. This possibility entails a significant increase in complexity and computing time. In principle, the dose should be computed for all the possible orientation of the beam accelerator, with continuity.

Material and Methods: To reach our goals, we train a generative Neural Network architecture, a Variational AutoEncoder (VAE) with both Graph encoding and decoding, to reproduce the dose distribution obtained with a Geant4 full MC simulation.

The peculiarity of our Network is in the pooling operations, which are done imposing on the graphs smaller adjacency matrices obtained recursively disconnecting nearest neighbours nodes and connecting 2nd nearest neighbours. Alternating GraphConv layers [2] and pooling operations, we encode the original graphs in smaller dimensional representations.

The decoding uses the same graph representations employed for the encoding, but in reverse order. Nodes from a lower dimensional representation are put in a bigger graph structure in which all other nodes are initialized to zero, then a Graph Convolution is executed. The process repeats until the original graph structure of data is recovered.

Once trained, we use the decoder part of the VAE to generate new dose distributions, sampling the latent space as a function of the energy and beam orientation.

Preliminary results: We have already obtained encouraging results regarding the generation of the dose distribution in a homogeneous material conditioned to the beam energy.

We are currently working on the generalisation to inhomogeneous materials. We will test the Network's ability to handle such increase in the task's complexity and, eventually, try different architectures and Graph layers.

1 E. Schüler et al., Very high-energy electron (VHEE) beams in radiation therapy; Treatment plan comparison between VHEE, VMAT, and PPBS Med. Phys. 44 (6), June 2017

[2] Morris, Christopher, et al. "Weisfeiler and leman go neural: Higher-order graph neural networks."Proceedings of the AAAI conference on artificial intelligence. Vol. 33. No. 01. 2019.