Version: 1.00

Enabling AI for High Energy Physics Experiment and Theory

Editors:

Alex Tapper, Biagio Lucini, Davide Costanzo, Monica D'Onofrio, Rhiannon Jones, Tim Scanlon

Contributors:

Mark Hodgkinson, Gert Aarts, Henning Flaecher, Dominic Brailsford+, Sam Van Stroud*, David Sankey++, James Frost, Mark Wilkinson, Eduardo Rodrigues+, Andy Blake+, Andy Chappell, Gavin Hesketh, Gurtej Kanwar, Benedikt Maier*, Fabrizio Salvatore, Karolos Potamianos, Sudarshan Paramesvaran++*

**Editor of Hardware, Software and AI-Ops SubSection*
+Editor Skills/Training, Capacity Building and Knowledge Transfer
***Editor of Industry and Wider Engagement SubSection*
++Editor of Fast-ML SubSection

Executive Summary: Over the past three decades, High Energy Physics (HEP) has successfully leveraged Artificial Intelligence (AI) across numerous aspects, including data analysis, theory calculations, detector calibration/monitoring, and real-time data selection. AI's importance in HEP is only set to grow further, and it is essential to the future of the field that it fully exploits its potential, but the UK community currently lacks a strategic and coordinated approach to achieve this goal. Establishing a UK AI HEP framework is essential to overcome common barriers and challenges while harnessing opportunities in areas such as hardware, software, AI-ops, skills/training, knowledge exchange, capacity building, industry engagement, and fast-AI. To address this need, a community workshop was held in October 2024, and this document details its findings, including a concise summary of recommendations. **Appendix A** contains a more detailed workshop summary, and **Appendix B** contains case-studies from the application of AI in HEP.

Findings

1. *Importance of AI:* AI is crucial to HEP, with its impact expected to transform almost every area of the field and enable breakthroughs that would be impossible with traditional methods; without swift adaptation and a strategic focus on AI, the UK risks falling behind other nations in this rapidly evolving domain.
2. *Dangers of falling behind:* If UK HEP does not embrace AI, it risks losing global competitiveness, missing key scientific opportunities, struggling to attract top talent, inefficiently utilising resources, weakening its role in international collaborations, and limiting its wider socio-economic impact. Adopting AI is essential to maintain leadership and capitalize on emerging opportunities in the field.
3. *Funding*: Dedicated funding for AI in HEP is infrequent, difficult to access, and lacks the stability needed for sustained research and implementation. Similar issues exist in the wider software and computing realm, with long-term support essential to fully exploit opportunities, as demonstrated by effective initiatives like Swift-HEP
4. *UK AI Forum*: There is a lack of a regular UK forum to discuss common barriers, challenges, and opportunities in AI across HEP.
5. *International Initiatives*: There are several international initiatives (for instance EuCAIF, FAIR for AI) also developing AI strategies, which are not widely known in the UK HEP community and sometimes lack sufficient UK engagement.
6. *Hardware, Software, and AI-Ops*:
    a. Access to GPU resources varies widely across institutes and users are often unaware of how they can access such resources, especially at the national and international level.
    b. Different experiments experience very similar technical issues, for instance those related to training, inference, and model stability against changing versions of AI software tools, but lack mechanisms to exchange expertise.
7. *Skills/Training, Capacity Building, and Knowledge Transfer*:
    a. Although there have been some very successful AI-training initiatives in the UK, for instance, STFC's/UKRI's Centres for Doctoral Training (CDTs) which have been hugely successful in imbedding expert AI practitioners into the field with revolutionary outcomes (see **Appendix B**), there is a lack of a coherent and concrete approach to such training across the UK. More can be done to expand the benefits of and learn from successful programmes, and to make a more enriched training offering across the field.

   b. An increasing number of physics PhD applicants are seeking research opportunities that incorporate AI, with the intention of applying these skills in either other research areas or industry throughout their careers. As such, a greater focus on AI training, challenges and opportunities, can significantly enhance PhD recruitment outcomes, attracting more talent to the field. This also helps embed expert practitioners of AI, with HEP domain expertise, into the field, which is also very effective for transferring foundational AI knowledge. The strong recruitment evident in the CDTs, which bring complementary new talent into the field, are a good example of the effectiveness of this approach.
   c. Some areas of HEP are enriched with AI expertise, both foundational (Hartree, SciML) and applied (various experiments), while others would benefit from more dedicated expertise. The field stands to gain significantly from ensuring more effective knowledge transfer within the field. There are also gains to be realised from engaging with foundational AI experts, especially for innovative and non-standard approaches.
   d. There is strong competition across sectors for AI experts, particularly those with the skill set developed in HEP. New initiatives, such as joint roles between industry and HEP, have proven successful in attracting these researchers and encouraging longer-term retention in the field.
8. *Industry and Wider Engagement*:
   a. Industry plays a leading role in the development of cutting-edge AI tools and techniques. The HEP community can benefit from aligning with industry tools to leverage developments and lower the threshold for expertise exchange.
   b. In some programmes there has been a successful track record of industry partnerships and placements, undertaken in a wide array of sectors, focussing on skills, techniques and hardware. This has created extra knowledge exchange opportunities and access to additional funding streams. However, this is not widespread across the field and there is the potential for significant gains from enhancing the way we work, improving our effectiveness and delivery of outputs.
   c. There is a strong demand in industry for the development of talent within HEP, which has enabled the wider socioeconomic impact of HEP to be enhanced, providing another strong narrative for funding of HEP research.
9. *Fast-ML*: applications in which AI inference is required on microsecond timescales or below, usually necessitating the use of specialized computing platforms, and often in the context of real-time data-selection systems.
   a. The UK HEP community has developed world-leading expertise in high-speed electronics and data processing algorithms, including AI algorithms, as well as in real-time data-selection systems (triggers)
   b. UK activities are in many cases independent, with limited collaboration and sharing of expertise.
   c. Science-led, open-source tools, such as hls4ml and conifer, which enable fast AI, primarily on FPGAs, have emerged, with early contributions from the UK.
   d. Working together with leading FPGA vendors, the UK HEP community can be influential in spreading the benefits of microsecond-scale AI to applications in industry where it is currently generally not widely developed/supported.

Recommendations

1. It is essential for the UK to adopt a strategic approach to AI for HEP, including sustained funding, skills development, talent retention, knowledge transfer, industry collaboration, and participation in global AI-initiatives.
2. Provision of dedicated, stable and long-term funding to ensure the UK is at the forefront of harnessing AI, whilst increasing the socioeconomic impact of the field.
3. Establish a regular forum for discussing common barriers, challenges, and opportunities in AI across the HEP field. An in-person meeting twice a year, circulating around different venues, was decided as an appropriate forum. Working groups focused on specific areas can meet more frequently as appropriate.
4. Increase engagement of UK-based researchers in AI strategy initiatives beyond-HEP, ensuring the UK's interests are represented and aligned with the outcomes.
5. Centrally maintain and advertise details of AI hardware and development environments available to HEP researchers.
6. Share expertise, document best practices and investigate shared frameworks/tools to aid knowledge exchange within the UK.
7. Develop a coherent approach to AI training in HEP in the UK, for researchers of all career stages, including:
    a. More stable and focussed delivery of STFC's data-science summer school.
    b. Build upon successful HEP AI training programmes, such as the CDTs, to ensure successful elements are more widely adopted and benefits widely felt.
    c. Initiate focused AI software carpentry-style workshops.
    d. Encourage more expert-exchange placement, with industry and within HEP.
8. Enhance AI knowledge exchange in the UK by mapping expertise, creating common forums/software, enabling expert practitioners to work across experiments, and building strong cross-experimental cohorts.
9. To better exploit their deep foundational expertise, build stronger links between HEP researchers and STFC's national facilities such as Hartree Centre and the SciML.
10. Where most essential, strengthen links to foundational AI expertise beyond HEP, by identifying and building links with motivated foundational researchers.
11. Ensure HEP is sufficiently attractive to retain/attract AI experts, in terms of job progression, security and making an attractive offering. Explore innovative offerings, such as joint industry and research roles.
12. Identify applications or projects (individuals, collaborative or consortia) that could be quickly initiated for future wider funding calls.
13. Exploit more efficient industry tools to maximally leverage developments and position for expertise exchange.
14. Maximise the wider socio-economic impact of HEP, enabling the field to tap into innovation funds and to build a stronger narrative for long-term investment from both government and industry.
15. Build on the UK HEP community's expertise in high-speed electronics and data processing algorithms to further develop Fast-ML activities, ensuring strong interconnections through a UK forum for groups of experts or interested parties.
16. Continue to develop leading-edge solutions in μs-scale AI, growing industry engagement, and making outcomes available to other research areas and beyond.
17. Commit to further contributions to development, support and training for open-source toolsets to enable fast AI, primarily on FPGA, for instance hls4ml and conifer.

Appendix A - Detailed Summary of AI Workshop

Introduction

Editors: Alex, Biagio, Davide, Monica, Rhiannon, Tim

1 HEP has a long and successful history of exploiting AI to further our research outputs, stretching over the past four decades. It is already used in many experiments across a wide array of areas, for instance data analysis, theory calculations, detector calibration/monitoring, and real-time data selection, in addition to efforts to accelerate Lattice field theory calculations. A dedicated community workshop was organised in October 2024 to evaluate the status of AI in the field of HEP theory and experiment. Based upon this workshop, a series of recommendations have been formulated and reported in this document.

2 If UK HEP does not embrace AI, it risks losing global competitiveness, missing key scientific opportunities, struggling to attract top talent, inefficiently utilising resources, weakening its role in international collaborations, and limiting the economic and societal benefits of technology transfer. Adopting AI is essential to maintain leadership and capitalize on emerging opportunities in the field.

3 AI is and will continue to be vital to many areas of experimental and theoretical particle physics. Nonetheless the HEP community lacks a common strategy of how to fully exploit it. There is a need to formulate a coherent strategy to enable AI in HEP by overcoming barriers, challenges and issues, whilst building upon opportunities. In the following, ambitions and recommendations focus on the areas of (i) hardware, software and AI-ops; (ii) skills/training, knowledge exchange and capacity building; (iii) industry and wider engagement; and (iv) fast-ML.

4 The need was identified for a regular forum to discuss common barriers, challenges and opportunities in AI across HEP. An in person meeting twice a year, circulating around different venues, was decided as an appropriate forum at this point. Working groups focussed on specific areas can meet more frequently as appropriate, and report in the bi-annual meeting.

5 There are other initiatives ongoing internationally to develop AI strategies, not widely known in this community. It is important that more UK-based researchers engage in these working groups, to ensure the UK community is kept updated on these efforts, represent the UK's interests and ensure UK efforts are aligned with these initiatives.

6 Dedicated funding opportunities to support AI in HEP are infrequent and difficult to access. Where they do exist, they usually do not provide the required stability for a sustained and fully developed programme. Sustainable and long-term funding to ensure the community is able to fully exploit the opportunity afforded by AI is essential.

Hardware, Software and AI-Ops
Editors: Alex, Andrew, Benedikt, Leigh, Luke, Mark, Sam

1 Despite there being numerous facilities available for training expensive AI models, users are often unaware of these facilities or how to access them, especially at the national and international level. We recommend identifying, documenting and maintaining details of different resources that UK particle physicists have access to in order to train ML models.

2 Different experiments experience similar technical issues, for instance those related to training, inference and model stability against changing versions of AI software tools. A common forum would allow such issues to be discussed amongst UK particle physicists, to converge on and document common best practices, initially, and in the longer term to investigate potential common frameworks and tools.

3 In the near future, there will be funding calls to access large scale GPU resources. We recommend that the UK community identify several high priority applications or projects that could be quickly prepared for such calls.

Skills/Training, Capacity Building and Knowledge Transfer
Editors: Andrew, Dominic, Eduardo, Tej, Tim

1 Although there are various AI training initiatives available, nationally and internationally, there is currently a lack of a coherent and concrete approach to AI training in HEP across the UK. While we have some examples of successful AI skills training in HEP, such as via the Centres for Doctoral Training, we need to expand these initiatives and develop additional training offerings accessible across the community and all career stages. This can include a more stable summer school offering, UK-based software carpentry style workshops (making use of, contributing to, and helping create HSF training in this area), and enhanced training for PhD students. A greater focus on AI training and expertise can significantly enhance PhD recruitment outcomes and attract talent to the field. This also helps embed expert AI practitioners, who also have HEP domain expertise, into the field, which is effective for foundational AI knowledge transfer into HEP.

2 Within HEP we have some research areas that are enriched with AI expertise and other areas which could significantly benefit from such expertise. We should ensure that there are mechanisms to transfer this expertise across the HEP domain. There are several different options, which are complementary and should be explored, including AI fora, development of HEP-specific software, mapping of community expertise and building stronger links across the community. The field stands to gain significantly from ensuring more effective knowledge transfer within the field. There are also gains to be realised from engaging with foundational AI experts, although these are more difficult to realise. However, to maintain a conversation with such experts and ensure exposure to the latest developments, HEP researchers should be encouraged to attend and present at wider-AI conferences/workshops.

3 There is very strong competition for researchers who are expert practitioners in the application of AI. We need to ensure that our field is sufficiently attractive that enough stay in the field. Options are needed for career progression, job security, ability to stay at the cutting-edge (training, hardware, software), and ensuring an engaging work environment. It

should be noted that career progression and job security are significant challenges for technically focussed researchers in academia, which demands more unique solutions to the issue. There are good examples of positions being created which offer a dual role working on distinct projects in both industry and HEP, which have been successful in retaining applied AI experts in the field. Such roles are part-funded by a partner who wants to engage the researcher's expertise, the researcher gets to enhance their training, career prospects and solve interesting real-world challenges, whilst also working on applying cutting-edge AI to further fundamental physics. The knowledge-exchange, talent retention, industry income and training opportunities, can lead to benefits for all parties.

4 There is already extensive AI expertise within the STFC community, for instance the Hartree Centre and the SciML groups, who both have significant foundational AI expertise. However, HEP has only weak links with these groups, especially the SciML group. We recommend that stronger links are built between these STFC areas, with a first step to launch discussions to identify the best ways to engage, enabling HEP to better exploit foundational expertise where needed, at a low cost.

Industry and Wider Engagement

Editor: *Gabriel*

1 Industry plays a leading role in the development of cutting-edge AI tools and techniques. The HEP community should ensure to fully engage with industry tools to maximally leverage developments and position for expertise exchange.

2 There has been a successful track record of industry partnerships and placements, undertaken in a wide array of sectors, focussing on skills, techniques and hardware. This has created extra knowledge exchange, additional funding streams and opportunities.

3 There is a strong demand in industry for the talent developed within HEP. This can help to maximise wider socio-economic impact, enabling the field to tap into innovation funds and to build a stronger narrative for long-term investment from both government and industry.

4 There can also be a significant gain from interacting with industry in terms of enhancing the way we work in research, improving our effectiveness and delivery of outputs.

Fast-ML

Editors: Alex, Alex, David, Sudarshan

1 Fast-ML consists of applications in which AI inference is required on microsecond timescales or below, usually necessitating the use of specialised computing platforms and often in the context of real-time data-selection systems.

2 Over the last two decades the UK HEP community has developed world-leading expertise in high-speed electronics and data processing algorithms, including AI, primarily for detector trigger and data acquisition systems. The UK should build upon this knowledge and expertise, to exploit opportunities in this area.

3 UK activities are in many cases independent, with limited collaboration and sharing of expertise. The UK should build stronger interconnections, via common forums for groups with common expertise or interest, whilst also increasing engagement with the nascent international community in this area.

4 Microsecond scale AI is not a widely developed and supported area in industry. The UK should continue to develop leading-edge solutions in this area, and grow engagement with industry and make the outcomes available to other scientific areas and beyond.

5 Science-led, open-source tools, such as hls4ml and conifer, to enable fast AI, primarily on FPGA, have emerged, with early contributions from the UK. The UK should commit to making further contributions to development, support and training for these toolsets.

Appendix B - Case Studies (work in progress)

Particle physics has a strong track record of exploiting AI across all aspects of experiment, including: physics object reconstruction, identification and calibration; real-time event selection; uncertainty evaluation; event selection; background modelling; and documentation. This work has been ongoing for over three decades and is continually adapting to exploit state-of-the-art AI solutions. The next generation of experiments will take data until the 2040s, and are prime examples of the need for long-term retention of AI skills and leadership, given their likely prevalence on current and these future experiments. AI is also a vital tool in particle physics theory and phenomenology, where the UK is also an international leader.

The case-studies below highlight UK intellectual leadership in international projects and demonstrate the impact of AI in High Energy Physics, which now commonly underpins large portions of the physics programme. The effectiveness of the Centres of Doctoral Training programme in embedding AI expertise and providing transformation change, has also been clearly demonstrated.

ATLAS

The identification of jet flavour (Flavour Tagging, FTAG) is central to the vast majority of the physics programme at hadron colliders and is critical for some of the highest-profile analyses, such as the search for Higgs pair production. Despite having exploited cutting-edge AI in FTag for over three decades, we have recently been able to make further paradigm shifting improvements in performance by exploiting the latest developments. For instance, the recent development of Graph Neural Network (GNN) and Transformer based models at ATLAS, driven by expert AI practitioners trained in STFC's Centre for Doctoral Training in Data Intensive Science at UCL, has led to a ground-breaking improvement in performance. These models have dramatically enhanced the flavour-tagging rejection in all regions, by factors of 2-3. Some of the largest improvements have been realised when probing the highest energy jets, where the efficiency has improved by over 60% for a fixed background rejection. This algorithm has also come with a myriad of other benefits, including, enabling greater explainability, the extraction of more physics knowledge, a simplified algorithmic approach, and dramatic gains in other areas. The impact of this work is demonstrated by the fact it has resulted in both an ATLAS Thesis Award Prize and 3 x ATLAS Outstanding Achievement Awards for members of STFC's CDT at UCL. This approach, aided by the versatile software package developed by the time, has also now been propagated to enhance numerous other areas on ATLAS including: primary vertex reconstruction, long-lived particles identification, tau-identification, H->bb identification, jet calibration and track reconstruction, demonstrating the wide ranging impact of this work. Further novel developments based upon image-segmentation techniques from industry are already promising further large gains in FTag and the ever challenging area of charged particle track reconstruction in the future. This all demonstrates the significant gains made possible by training students who are both HEP domain experts and expert practitioners in AI, who are equally comfortable in both areas and act as a knowledge conduit between the two areas.

The STFC LIV.INNO CDT has also trained diverse cohorts of students in data intensive science, addressing the data challenges presented in astronomy, nuclear, theoretical and particle physics, accelerator science, mathematical and computer science. It followed the successful LIV.DAT CDT programme, also relevant for AI applications in HEP. In ATLAS, students supported by the Liverpool CDTs have developed NN for searches of new physics, i.e. for hidden sectors such as the dark photons search, where a convolutional NN exploiting the 3D representation of calorimeter energy deposits associated with jets has developed. Generated additional activities thanks to the CDT programmes support include multidisciplinary consortia (EU funded CHIST-ERA MUCCA on explainable AI) and collaborations in health science (e.g. use of AI models for predictions in Age-related Macular Degeneration and Diabetic Retinopathy, thesis in progress with ILCaMS at Liverpool).

A large number of ATLAS-UK physicists from a range of groups have also been heavily involved in the crucial di-Higgs analyses, where cutting-edge AI models have been deployed in various areas to enhance these flagship analyses. This has included AI models to enhance H->bb identification (see above), event classification and background predictions. In particular, normalising flow techniques have been deployed to enhance the background determination in the HH->bbbb channel by UCL DIS CDT PhD students. This has led to a more robust background prediction, with reduced systematic uncertainties, and has paved the way for an AI-based classifier for signal-to-bkg discrimination. This was previously not possible/useful since any improvements from using an AI classifier were washed out by the increased systematic uncertainties. The new approach has both simplified the analysis and led to a 30% improvement in expected sensitivity. Such AI-driven improvements are playing a critical role in the LHC's ability to identify di-Higgs production, which will be a ground-breaking result and one of the legacies of the LHC.

CMS

The upgrade of the LHC at CERN will result in very challenging conditions for detectors with up to 200 simultaneous proton collisions 40 million times a second. Determining which of these collisions are interesting enough to keep to later analysis in a few millionths of a second requires state of the art technology. UK CMS members have developed an AI algorithm, based on convolutional and dense Neural Networks, to run in a dedicated chip sifting through the trajectories of the particles produced in these collisions in a few hundred billionths of a second and deciding which to keep. AI is much more efficient than existing algorithms for this task, which is a crucial pillar of the entire physics programme. The impact of this work is also demonstrated by the fact it has resulted in prestigious CMS thesis awards.

CMS UK members also developed a fast Convolutional Neural Network algorithm for topological searches at High Luminosity LHC. The algorithm has the potential to improve signal efficiency for flagship measurements such as di-Higgs to bottom quark decays. The algorithm was implemented on an FPGA, for use in the first level trigger system and conformed to latency and resource limitations. The AI showed an alternative method for triggering was possible which can be used in conjunction with classical techniques to enhance fast data selection and extract more information from the LHC datasets. The algorithm implementation had led to a collaboration with the NHS on how CNNs can be used in medical imaging analysis. This work again demonstrates the impact of the STFC CDT programme, through the STFC Data Intensive - Centre for Doctoral Training Cardiff, Bristol and Swansea.

A key example of Neural Network usage in CMS data analysis is in a search for new particles decaying to pairs of Higgs bosons in the di-photon+di-tau final state performed to investigate several related excesses present in earlier datasets. A parametric neural network was trained to search optimally across a wide range of parameter space. This technology allows a single network to adapt to very different kinds of signals with different kinematics. Without this AI technique, the sensitivity would be limited to a very specific set of signal models.

As a second example, CMS measured Higgs cross-sections in the diphoton final state at 13.6 TeV with the latest LHC data This analysis is one of the main channels used to characterise Higgs boson production. A new method for calibration of the showering simulation was developed using normalising flows, which allows for corrections of the simulation in a very high dimensional space. This calibration allowed us to significantly reduce the main experimental systematic uncertainty, related to the showering model, which would otherwise limit the sensitivity of these measurements. This work included contributions from students on a AI dedicated masters programme and CDT at Imperial.

LHCb

LHCb physicsts have been using AI for over a decade and it is an essential part of the experiment, enabling the physics programme.This is particularly true in the trigger system, the very first stage of the selection of the data for analysis, which has been relying on AI (with Boosted Decision Trees) from early on. The full detector is presently read-out at a rate of 32 Terabits per second, the highest ever achieved. Already in the first stages of the (trigger) selection of what eventually will make it for analysis, AI algorithms efficiently disentangle signals from uninteresting events, with signal production rates being as high as of the order of the million per second.

More data than ever also requires more simulation samples, hence a need for (ultra) fast simulations. ML-based parameterisations for simulation result in speed-up factors of two orders of magnitude, a must to be able to produce simulation samples commensurate with the needs of the data samples being analysed.

AI/ML is used across the entire LHCb physics research programme, from the trigger system, to particle identification, B-meson flavour tagging, data quality monitoring, simulation and physics analysis. Examples include: monotonic (Lipschitz) neural networks in the trigger; MLPs, BDTs and others for physics analysis; GANs for simulation; Deep Sets for inclusive flavour tagging of B mesons; GNNs for track finding; DL based algorithm for full-event interpretation. Many new approaches and methods are in production, others under investigation.

DUNE/SBND

The DUNE-UK Reconstruction Software and Distributed Computing project is an example of STFC-funded AI work in the neutrino sector, with almost half of the project related to AI.

AI-driven reconstruction enhancements show improvements to prospective physics sensitivity expectations for the DUNE physics programme. The main AI use cases in LArTPCs are related to pattern recognition, particle identification and neutrino flavour characterisation and enhance the ability to efficiently select signal over the background.

UK physicists have made major contributions, including the development of a flavour-tagging CNN, enhancing sensitivity to CP violation in the lepton sector and mass hierarchy determination. This demonstrated for the first time that DUNE can achieve its primary physics goals as laid out in the Conceptual Design Report. Other CNN based tools for separating track and shower-like energy deposits in LArTPC were developed in the UK and used in all of the ProtoDUNE analyses. These developments contribute to the ProtoDUNE physics goals and serve as vital prototyping for the final DUNE Far Detectors.

Another example is the development of AI-based neutrino vertex reconstruction in the Pandora package for MicroBooNE and DUNE using U-ResNet semantic segmentation. This will benefit many physics use cases, including atmospheric neutrino searches, long baseline neutrino oscillations and proton decay searches.

In addition, there is UK-led work SBND semantic segmentation for cosmic-ray background removal in neutrino interaction searches and the UK is involved in a project to use FastML for supernova shower neutrinos to predict the timing and direction of their occurrence.

Lattice Gauge Theories

Lattice Gauge Theories provide a robust calculational framework from first principles for strongly interacting gauge theories. Benefiting from 50 years of development, they are now a reliable tool for evaluating observables in QCD and in theories of strong interactions beyond the Standard Model. Current applications to QCD include isospin breaking effects, heavy meson decays, QCD in extreme conditions, and the calculation of the hadronic contribution to the muon g-2, while beyond the Standard Model applications include the determination of spectra and decay constants that can inform future experimental searches at particle accelerator facilities.

At the heart of the lattice approach to gauge theories are Monte Carlo calculations running on state-of-the-art massively parallel computer architectures. These highly demanding numerical calculations on flagship supercomputers are paramount to reaching the needed accuracy to compare with experimental results. The workflow used in lattice calculations can be broken into two main steps: generating the relevant field configurations and calculating the observables over these configurations. The numerical demands of the generation of the configurations are due to the need to extrapolate to infinite volume and zero lattice spacing the results of the calculations, while the intensity of the measurements is due to statistical noise arising from the quantum fluctuations that characterise the system.

AI methods have been employed to reduce the demands of both components. For the generation of the configurations, the high-dimensional probability distribution of the system has been linked to simpler functional forms obtained by continuous evolution through learned normalising flows or diffusion models. The use of AI has been crucial to continuously deform the distributions from trivial priors to highly correlated target distributions and to determine an optimised evolution dynamics. To implement gauge symmetries efficiently, the development of gauge-equivariant formulations has been shown to be essential. Investigations on theories in fewer than four spacetime dimensions show that these methods enable us to significantly reduce the computational time needed for a fixed accuracy, sometimes by several orders of magnitude. Scaling these methods to theories in four spacetime dimensions is the object of current investigations. In evaluating observables, significant progress has been made using convolutional neural networks to construct improved estimators or to relate scaling across systems of different sizes at phase transition points, again reducing significantly the time needed for obtaining a given accuracy. The use of AI has been crucial in solving long-standing problems such as the construction of order parameters in the absence of explicit symmetry-breaking terms in the Hamiltonian and the definition of a physically motivated inversion of the renormalisation group flow. As in the case of applications of AI to generate configurations, current research is focussed on optimising the proposed methods for efficient applications to mainstream lattice calculations. Besides the above applications, which fall under AI4Science, there is an effort in the opposite direction, Science4AI, to robustly understand the process of learning. A particularly successful direction employs Random Matrix Theory, a method well known in Lattice gauge theory, to investigate learning as a stochastic process in a high-dimensional matrix space This analysis has provided invaluable insights into what can be learned by given sets of training data.

Phenomenology

By embedding the principles of equivariance in the rapidity-azimuth plane and infrared and collinear safety (IRC) into a graph neural network, performance increases in the isolation of hadronic jets arising in strongly coupled dark sector scenarios was demonstrated. The algorithm also exhibits an enhanced optimisation capability compared to approaches that do not exploit symmetry, i.e., that are physics-opaque. Its potential as a robust tool for advancing search strategies for, e.g., semi-visible jets will have an impact on the future roadmap of particle physics that will see more searches for sparse and soft hadronic signatures.

Development of a neural network based model to emulate matrix elements improves on existing methods by taking advantage of the known factorisation properties of matrix elements. Thereby, one can control the behaviour of simulated matrix elements when extrapolating into more singular regions than the ones used for training the neural network. The performance of the model is demonstrated in the case of leading-order jet production in e^+e^-collisions with up to five jets. The model can reproduce the matrix elements with errors below the one-percent level on the phase-space covered during fitting and testing, and is robust in the extrapolation to the parts of the phase-space with more singular matrix elements. This was expanded to include coloured initial states and accelerate unweighted event generation of Z+jets events.

With the advent of widely available noisy intermediate-scale quantum computers (NISQ), there is increasing interest in quantum algorithms applied to high-energy physics problems. A method for anomaly detection based on autoencoders implemented via variational quantum circuits has been developed that outperforms its classical counterpart for small datasets, demonstrating quantum advantage in regimes of low sample size. Such highly efficient algorithms of unsupervised learning could soon run on actual NISQs to explore model-agnostic, data-driven approaches in the search for new physics.

Parton distribution functions (PDFs), which describe the probability of finding a specific type of parton (quark or gluon) carrying a certain fraction of the proton's momentum inside a proton, are crucial for making accurate predictions in collider phenomenology. An open source framework made available by the NNPDF collaborationexploits state-of-the-art developments in AI to realise a comprehensive determination of the proton structure from a wealth of experimental data. In contrast to other fitting approaches, where the PDF shape is parametrised in terms of relatively simple functional forms more or less inspired by QCD models, artificial neural networks are used as unbiased interpolants. The availability of this framework as open source will encourage the broader high-energy and nuclear physics communities to deploy machine learning methods in the context of PDF studies.