



ROADMAP

Roadmap on fast machine learning for science

OPEN ACCESS

RECEIVED

20 May 2025

REVISED

19 December 2025

ACCEPTED FOR PUBLICATION

19 February 2026










PUBLISHED

17 March 2026

Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](https://creativecommons.org/licenses/by/4.0/).

Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.



Sioni Summers^{1,11,*} , Alex Tapper^{2,11,*} , Thea Klæboe Årrestad³ , Chen Qin⁴, Karin Rathsmann⁵ , Matthew Streeter⁶ , Charlotte Palmer⁶ , Jonathan Citrin⁷, Changgang Zheng⁸ , Noa Zilberman⁸ , Alexander Titterton⁹ and Tobias Becker¹⁰ 

- ¹ CERN, 1, Esplanade des Particules, Meyrin, Switzerland
 - ² Department of Physics, Imperial College London, London, United Kingdom
 - ³ ETH Zürich, Otto-Stern-Weg 5, Zürich, Switzerland
 - ⁴ I-X and Department of Electrical and Electronic Engineering, Imperial College London, London, United Kingdom
 - ⁵ European Spallation Source ERIC, Lund, Sweden
 - ⁶ School of Mathematics and Physics, Queen's University Belfast, Belfast, United Kingdom
 - ⁷ Google DeepMind, London, United Kingdom
 - ⁸ Department of Engineering Science, University of Oxford, Oxford, United Kingdom
 - ⁹ Graphcore, 11-19 Wine Street, Bristol BS1 2PH, United Kingdom
 - ¹⁰ Maxeler Technologies, a Groq Company, 3 Hammersmith Grove, London W6 0ND, United Kingdom
 - ¹¹ Guest editors of the Roadmap.
- * Authors to whom any correspondence should be addressed.

E-mail: sioni@cern.ch and a.tapper@imperial.ac.uk

Keywords: machine learning for science, big data, particle physics, codesign, coprocessors, heterogeneous computing, fast machine learning, large language models

Abstract

The need for microsecond speed machine learning (ML) inference for particle physics experiments has emerged in recent years, in particular for the forthcoming upgrades to the experiments at the Large Hadron Collider at CERN. A community has grown around the need to develop the custom hardware platforms and tools required. The material presented in this report is drawn from the latest workshop held by the fast ML for science community and comprises of a collection of perspectives on the status of fast ML in different scientific domains, and the supporting technology.

Contents

1. Introduction	3
2. Fast ML at the Large Hadron Collider (LHC) experiments	4
3. Deep learning (DL) for fast MR imaging and analysis	7
4. Fast ML for accelerator controls	10
5. Fast ML for laser-plasma accelerators	13
6. Fast ML for fusion simulation, optimisation, and control	15
7. In-network ML: inference at the speed of data	18
8. Accelerating traditional HPC using AI: a selective overview	21
9. Inference speed is key to unleashing the potential of LLMs	24
Data availability statement	26
References	26

1. Introduction

*Sioni Summers*¹ and *Alex Tapper*²

¹ CERN, 1, Esplanade des Particules, Meyrin, Switzerland

² Department of Physics, Imperial College London, London, United Kingdom

E-mail: sioni@cern.ch and a.tapper@imperial.ac.uk

The fast machine learning (ML) for Science workshop was hosted by Imperial College London, from 25–28 September 2023, the fourth edition of the workshop.

The genesis of the workshop series has been the need for microsecond speed inference for the High-Luminosity LHC detectors, in particular in the hardware trigger systems of the ATLAS and CMS experiments. This level of speed requires non-standard and generally custom hardware platforms, which are traditionally very challenging to program. In addition, while ML is becoming widespread in society, this ultrafast niche is not well served by commercial tools. Consequently the field of particle physics has developed tools and techniques and a community of people in this area.

The workshop brought together almost 200 scientists and engineers in a hybrid format, to discuss the latest developments in fast ML. Participation from students, including some undergraduates, and early career researchers was a particular feature of the workshop, alongside representation from key industry partners. A strong aim of the conference was to engage scientific communities outside particle physics and develop areas where the tools and techniques from particle physics can be game-changing for other scientific fields.

The interdisciplinary nature of the workshop, including fields such as particle physics, free electron lasers (FELs), nuclear fusion, astrophysics, computer science and biology, made for a varied and interesting agenda, with much to discuss in coffee breaks and elsewhere. In addition to particle physics, the attendees heard talks on how fast ML is being harnessed to speed up identification of gravitational waves, to improve multi-messenger astronomy, how in FEL experiments, ML is needed to handle high data rates and due to the fast turn around of experiments, how flexible and generic ML features are necessary to allow for optimal use of the limited time at the facilities. Speakers covered medical uses for fast ML, through the need for fast image processing and data analysis for diagnosis and treatment and topics in biology searching for known and unknown features in large, heterogeneous datasets, for example common features in different types of tissues. The use of ML in control systems and simulations was discussed in the context of laser accelerators and in nuclear fusion experiments. Even theoretical physics featured through the application of ML to solve the electron wave equation in condensed matter, working towards a detailed and fundamental understanding of superconductivity.

Key industry partners, including AMD, Graphcore, Groq and Intel contributed to the workshop, discussing current and future generation hardware platforms and architectures, and running tutorials on using their development toolchains. Groq and Graphcore presented their latest dedicated chips for artificial intelligence (AI) applications, such as transformers and graph-based processing, and showed they have interesting applications to science problems. AMD and Intel highlighted the flexibility of their field programmable gate array (FPGA) platforms and explained how to optimise these for scientific ML applications.

The following articles are a collection of perspectives on the status of fast ML in different scientific domains, and the supporting technology.

2. Fast ML at the Large Hadron Collider (LHC) experiments

Thea Klæboe Årrestad¹ and Sioni Paris Summers²

¹ETH Zürich, Otto-Stern-Weg 5, Zürich, Switzerland

²CERN, 1, Esplanade des Particules, Meyrin, Switzerland

E-mail: thea.aarrestad@cern.ch and sioni.summers@cern.ch

Status

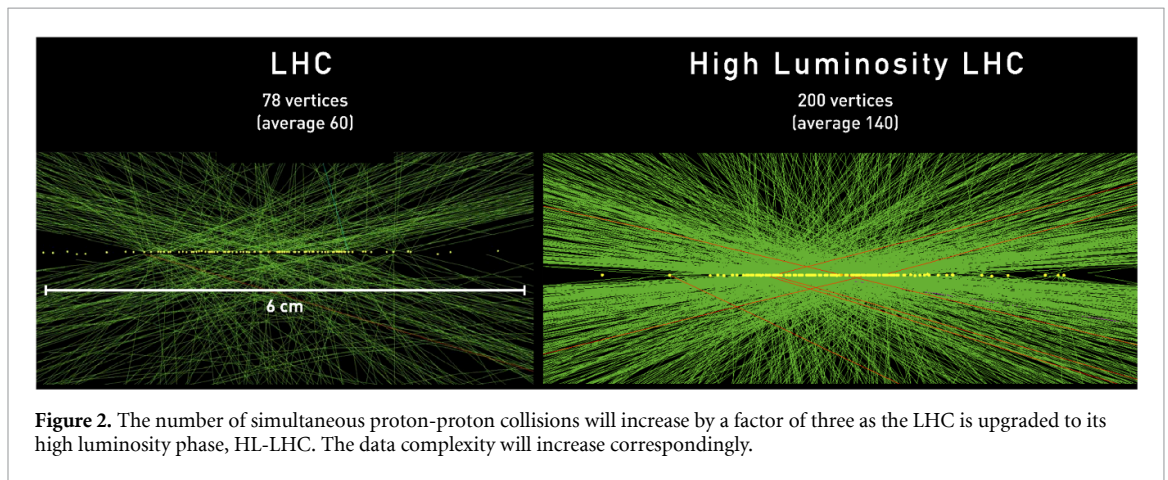
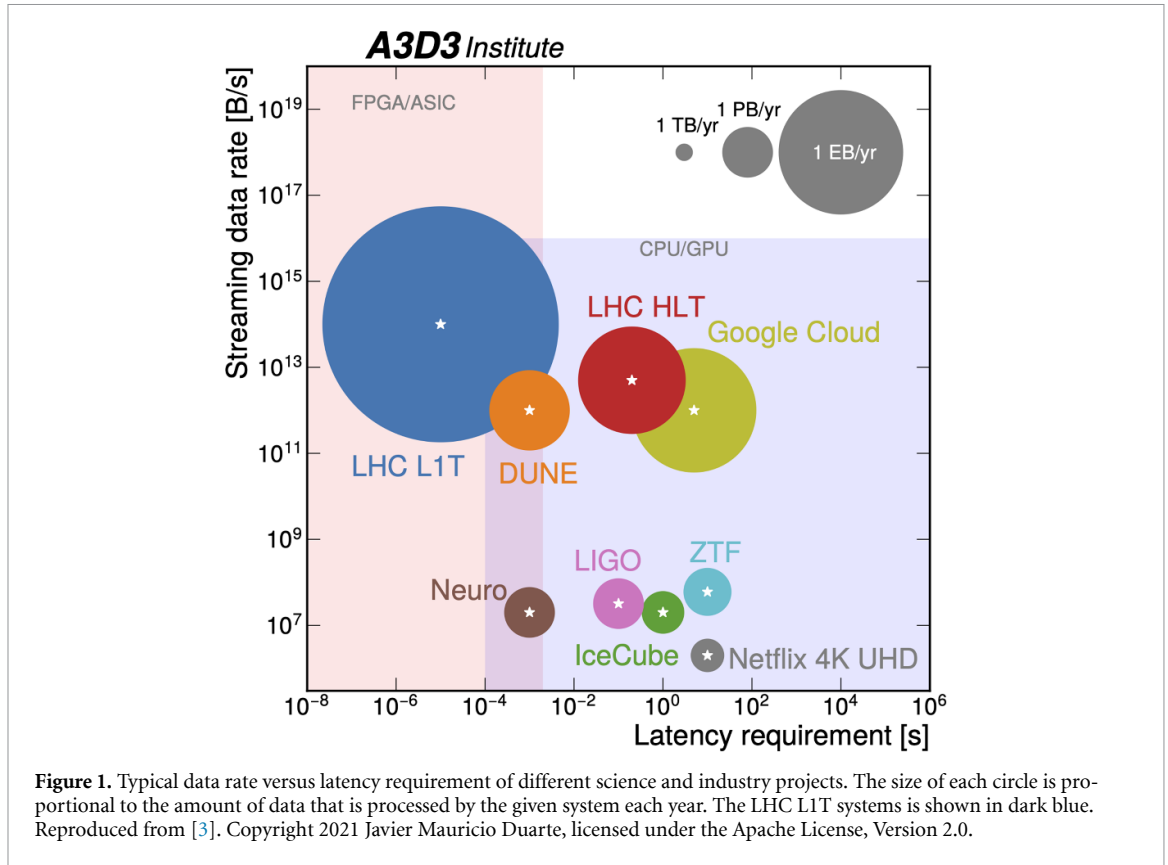
At the CERN LHC, bunches of trillions of protons are brought to collide in the centre of the four particle detectors around the LHC ring: ATLAS, CMS, ALICE and LHCb. This generates showers of new outgoing particles that produce signals in the particle detectors located around the interaction point. Due to the high frequency of collisions, every 25 ns, and large number of read-out channels, $O(10)$ million, an enormous amount of data is generated. For the two hermetic general-purpose experiments, CMS and ATLAS, this is $O(100)$ Tb s^{-1} . These signals are read out using custom electronics mounted on support structures inside the detectors. This large data rate can not be read out and stored as it would require a significant amount of electronics and power supply inside the detector, obstructing the path of generated particles and reducing the detector sensitivity. Rather, the data rate is reduced by a two-stage filtering system, the trigger. While data is buffered inside the detector on detector frontend electronics, a subset of the detector information is sent through high speed optical links to a radiation-shielded cavern located next to the detector. Here, the first stage of the event filtering system, the Level-1 trigger (L1T), is responsible for reducing the data rate by two orders of magnitude within the buffering time of $O(1)$ μs . Due to the strict latency constraints, this processing is done fully in firmware on $O(1000)$ FPGAs. The complexity of the data requires hundreds of reconstruction algorithms to run in parallel on several collision events simultaneously, which requires each algorithm to use minimal resources and time. The extreme latency and bandwidth requirements for the L1T system is unique and requires novel solutions in terms of IO and algorithms, as illustrated in figure 1. Due to their ability to efficiently and accurately process and parameterise high dimensional input, ML algorithms are increasingly being explored as faster and better replacements of the classical algorithms currently in use in the L1T. The extremely low latency and resource requirements for ML algorithms deployed in the L1T system has led to the development of dedicated software libraries that facilitates the training and translation of ML models into efficient FPGA firmware; *hls4ml* [1] for deep neural networks and *Conifer* [2] for decision trees. This has allowed for the first deep neural networks to be integrated into the L1T system, running inference within a few tens of nanoseconds.

Current and future challenges

With the rapid progress of *hls4ml* and *Conifer*, deep neural network and decision tree ensemble sub-microsecond inference in L1T systems is a reality. However, there are significant challenges ahead. In order to study increasingly rare physics processes in the LHC detectors, the LHC will be upgraded to its high luminosity phase (HL-LHC), allowing the experiments to collect a factor of ten more data. This entails a substantial trade-off: The number of concurrent proton collisions will triple, leading to a notably increased event complexity, as illustrated in figure 2. To cope with this, more granular detectors will be installed, demanding even more complex reconstruction algorithms due to the large increase in read-out channels. For the first time, the CMS L1T will also receive information from the inner tracking system. This will require the reconstruction of hundreds of charged particle tracks from thousands of detector hits, within a maximum latency of 5 μs . The input data rate to the L1T systems will increase by one order of magnitude, corresponding to around 5% of the total internet traffic [4]. This necessitates the design and implementation of significantly more advanced ML algorithms for the reconstruction, correction and selection of proton collision events.

Advances in science and technology to meet challenges

In order to meet the challenges imposed by HL-LHC, there is a strong ongoing R&D effort within the particle physics community towards the development of powerful neural networks running on specialized hardware like ASICs, FPGAs or AI accelerators. Owing to the point cloud characteristics of particle physics data, which involve sparse and unordered hits within irregular geometry detectors, graph neural networks (GNNs) have become the preferred architecture for most reconstruction tasks faced by the HL-LHC. Using an interaction network architecture, a type of GNN, ML-based charged particle tracking on FPGAs has been successfully demonstrated with $O(100)$ ns latency [5, 6]. For primary vertex



finding and association of tracks to the vertex, end-to-end ML solutions are also being explored [7]. Additionally, GNNs on FPGAs have been successfully demonstrated for energy reconstruction for highly granular calorimeters [8]. Direct access to the particles will for the first time enable the grouping of particles into jets at L1T [9]. This has led to a significant effort focused on using particle information to discern the flavour of the jet using ML, specifically with GNN-like architectures such as interaction networks [10] and transformers [11] or with set-based architectures [9]. In order to expand the reach of trigger systems beyond simple rule-based event selections, there has been a recent adoption of unsupervised learning for anomaly detection in the hardware trigger at both ATLAS and CMS [12, 13]. Besides the deep neural networks mentioned earlier, decision tree ensembles are also extensively used in the L1T owing to their superior performance on tabular data and their low latency and resource requirements. These are not employed for reconstruction tasks, typically focused on mapping an input set (like a group of detector hits) to another set (such as reconstructed tracks). Instead, they are used for tasks such as regression [14], classification [15], and anomaly detection. To account for fill-level changes in detector conditions, techniques like continual learning (CL) are being explored in order to maintain model stability in the resource-constrained, low-latency setting of the L1T [16].

Finally, exploration is also underway into the use of ML in the front end electronics of detectors. This includes on-sensor data filtering using neuromorphic computing-based spiking neural networks [17], on-pixel track seeding [18] and data compression in detector front end electronics [19, 20].

Concluding remarks

In conclusion, the integration of ML algorithms into the L1T system of the LHC experiments marks a significant advancement in particle physics data processing. ML, particularly deep neural networks and decision trees, has enabled ultra-low latency and efficient data handling, making real-time event selection and reconstruction feasible. This achievement is exemplified by the `hls4ml` and `Conifer` libraries, which facilitate the deployment of ML models on FPGAs with sub-microsecond inference times. However, as the LHC transitions to its HL-LHC to collect even more data, new challenges emerge. The increased number of proton collisions and the introduction of more granular detectors demand advanced ML algorithms to handle the higher event complexity. The community is responding with ongoing research and development, focusing on powerful neural networks, particularly GNNs, for tasks like particle tracking, energy reconstruction, and jet classification. Additionally, the adoption of ML extends beyond the trigger system and into the front-end electronics of detectors, enabling innovations such as neuromorphic computing-based data filtering and on-pixel track seeding. The pursuit of CL techniques ensures model stability in the resource-constrained, low-latency environment of the L1T. The fusion of ML and particle physics is not just a technological leap but a testament to the collaborative efforts of scientists and engineers working towards a common goal of gaining a better understanding of our universe while taking advantage of cutting-edge technologies to meet the evolving challenges of particle physics experiments at the HL-LHC.

Acknowledgments

T. Å. is supported by the Swiss National Science Foundation Grant No. PZ00P2_201594.

3. Deep learning (DL) for fast MR imaging and analysis

Chen Qin

I-X and Department of Electrical and Electronic Engineering, Imperial College London, London, United Kingdom

E-mail: c.qin15@imperial.ac.uk

Status

Magnetic resonance imaging (MRI) is a leading diagnostic modality for a wide range of exams and is one of the most useful and important imaging techniques in hospitals, due to its lack of ionising radiation and ability to probe various aspects of the physiology. MRI has been widely applied in neuroimaging, cardiovascular and musculoskeletal systems, and more than 50 000 scanners are estimated to be in use worldwide [21]. A typical MRI workflow ranges from patient scheduling and acquisition to image analysis and prognosis (figure 3) [22]. Conventionally, the MRI workflow can be very slow due to both the fundamentally limited acquisition speed of MRI and the laborious manual labelling for human analysis and interpretation. Recently, DL techniques have opened the possibility to accelerate this considerably, where they can impact all aspects of the MRI workflow. For instance, in cardiovascular MR (CMR) imaging, several acquisition-related aspects of CMR examination such as the automatic view planning can be automated or substantially shortened using DL [22, 23]. For MRI acquisition and reconstruction, DL has also been emerging as a popular and powerful alternative to the generic compressed sensing techniques. State-of-the-art DL-based methods have proposed to leverage routinely performed MR scans as training data to exploit prior knowledge as well as incorporating the physically acquired raw data in the reconstruction process [24, 25]. They can not only achieve high reconstruction quality but can also offer high efficiency in terms of reconstruction speed, which makes the clinical deployment feasible [22].


DL in MR image analysis and interpretation has also been widely investigated. DL in MRI segmentation has shown great potential in achieving human-level performance while significantly speeding up the annotation process [26]. For instance, a DL method [26] only takes 2.2 s to analyse images of a single subject at two time points of the cardiac cycle, whereas it typically takes a trained expert 20 min. DL also enables tasks that are impossible for a human to label, such as estimating dense correspondences between images or tracking the motion in sequences. They also offer much faster inference speed (more than 10x) compared to conventional optimisation-based approaches [27, 28]. According to the statistics in [29], automated image interpretation using DL has the potential to reduce the time radiologists spend on reviewing images by 20%.

Current and future challenges

Despite the remarkable performance and the tremendous promise that DL has shown in the field, there are several challenges and limitations that need to be mentioned. One of the major challenges is the limited data availability. DL methods are known to be data-hungry techniques, which rely heavily on the availability of training data. However, medical data is often expensive to obtain and the labelling of them can also be imperfect. Most current datasets for MRI images are typically in the range of a few hundreds, which may lead to a high risk of DL models being overtrained on the training data, and moreover, most likely there are more healthy subjects than diseased patients in the datasets, resulting in highly imbalanced classes. In addition, the sharing of data between institutions has also been challenging due to privacy regulations.

A relevant challenge that is related with the data limitations is the model bias and generalisation. Distribution shifts across sites, scanners, and acquisition protocols can significantly reduce model performance on unseen data [30]. Although there are some large population datasets available in recent years such as UK Biobank [31], there also exist biases towards groups with certain demographics. DL models trained on those can also lead to model bias towards certain groups of subjects and therefore result in concerns of fairness issues in the domain.

A further major challenge that currently limits the translational potential of such DL models in clinical practice is the trustworthiness [32], which encompasses model transparency, explainability, robustness, and safety. Although DL methods can achieve promising performance, there is often a lack of consideration of their reliability, e.g. they may show instabilities and lack of generalisability when encountered with perturbations and new domains [33]. On the other hand, the ‘black-box’ nature of DL approaches also often leads to the decision-making process being opaque, whereas it is crucial to make informed decisions in the safety-critical healthcare applications.



Computational scalability also represents a practical barrier. High-resolution volumetric imaging and multi-modal datasets demand substantial computational resources, making traditional pipelines slow and often impractical. Fast ML approaches can address this challenge, enabling efficient training and real-time inference, which are critical for deploying DL models in time-sensitive clinical workflows.

Advances in science and technology to meet challenges

Learning with limited or imperfect labels remains a challenge. Supervised learning with manual annotations is often constrained by that, and therefore methods that are beyond supervised learning are needed to tackle the challenge. For instance, self-supervised learning approaches [34] to learn useful representations without any human intervention can be investigated, such as via exploring the intrinsic structure within data itself or via contrastive learning to learn general features of datasets without labels. Unsupervised generative modelling is also another opportunity, where the generative power of DL models such as generative adversarial networks (GANs) or diffusion models [35] can be leveraged to generate synthetic labelled datasets for training. The self-supervised or unsupervised learnt representations are an effective way of exploiting unlabelled data and can be potentially leveraged to transfer knowledge to relevant new tasks where labelled data is sparse. Recent advancement of foundation models is a successful case of that, where models are trained in a self-supervised fashion on broad data and then be adapted to a wide range of use cases. There have also been explorations of foundation models in MRI [36], which mitigate the need of large data collection and labelling for each specific task. Such foundation models also have the potential to address the model generalisability and imbalanced dataset issues due to the large-scale datasets that are used [36]. To protect data privacy while sharing the multi-cohort knowledge, federated learning [37] can be leveraged to enable distributed learning across multiple institutions without central data sharing.

DL trustworthiness in medical imaging especially in MRI computing is also an important topic to tackle. It is necessary to equip the DL models with the ability of knowing what they do not know and inform the potential risks of failures. To achieve this, uncertainty modelling [38] can serve as a mechanism to communicate the knowledge boundary of DL systems and provide an additional reference for visualising and interpreting prediction reliability. The quantified uncertainty could also be leveraged to handle cases outside distribution and improve model confidence and robustness, potentially through test-time adaptation and uncertainty reduction techniques. Furthermore, DL trustworthiness in fast MR imaging and analysis could also be enhanced by improving model explainability either by design or using external techniques [39], such as with attention-based mechanisms or prototype learning methods. Causal analysis [40] is also a promising emerging field that can provide more insights within models and strengthen their clinical translation potential. Fast ML approaches can also complement these advances by making high-resolution, uncertainty-aware, and interpretable analyses computationally feasible. Efficient architectures, model compression, and hardware-optimized pipelines ensure that DL models can operate in real-time, enabling both clinical scalability and reliable deployment.

Concluding remarks

This section provides a brief introduction on the recent advancement of DL technologies for fast MR imaging and analysis. DL has the potential to improve the entire MR imaging workflow and has been shown to have greatly improved the efficiency of radiology workflow, allowing time for clinicians to focus more on patients' care. The section also identifies a few current and future research challenges that are limiting the DL development in medical imaging and hindering their clinical translation potential, which consist of issues related with data scarcity, model generalisability and model trustworthiness. Recent advancement in DL technologies that can meet the above challenges have also been discussed, including self-supervised/unsupervised learning approaches, foundation models and uncertainty modelling. The development of DL in medical imaging will continue to rise, and it would be expected to have a widespread impact on the future of fast MR imaging and analysis, ultimately bringing significant benefits to patients and the healthcare industry.

Acknowledgments

C. Qin acknowledges the support by the UK's Engineering and Physical Sciences Research Council (EPSRC) (Grants EP/X039277/1 and EP/Y002016/1).

4. Fast ML for accelerator controls

Karin Rathsmann

European Spallation Source ERIC, Lund, Sweden

E-mail: karin.rathsmann@ess.eu

Status

Accelerator-based research facilities include some of the world's most complex human-made systems. Operating a user facility, such as the European Spallation Source (ESS) in Lund (figure 4), is a challenge since scientific users will only have access to the facility during short periods in a tight schedule. This implies high availability demands, which for ESS is 95%. Other challenges and limiting factors for accelerator-based research are operational costs for electrical power, staff, maintenance and waste management (figure 5).

Given the recent progress in AI in general, there are high expectations that AI can address these challenges to get more science out of the accelerator-based facilities at reduced costs for the society and the planet.

To illustrate how existing ML algorithms can make a difference for automation and fault detection, we can imagine a subsystem with inputs, outputs, feedback loops and a goal for the subsystem to fulfil. If the output of the subsystem can be forecasted from the input using for example recurrent neural networks (RNN), then the subsystem could respond pre-emptively and accurately to changes in the input. If values of measured and predicted output would differ it could be an indication of an issue with the subsystem. This illustration also highlights similarities between control system networks and neural networks, which is not a coincidence since control systems and nervous systems share a similar task, namely, to regulate complex systems.

AI has been used for decades on the experimental side of accelerator-based facilities. Only recently it started to be applied to the accelerators themselves. The first workshop for AI in high-energy and nuclear physics was arranged in 1990 [41], compared to the first workshop for ML for accelerators in 2018 [42]. However, AI was the topic of 21 out of the 218 contributions at the international accelerator and experimental physics controls conference ICALEPCS in 2021 [43]. The applications of AI in these articles can mainly be categorized into fault detection, automation, optimisation, augmentation, as well as ML operation (MLOps) platforms for control systems.

Notable progress in fast ML for accelerator controls includes the development of a system to regulate magnetic power supplies at the Fermilab Booster [44]. This system uses reinforcement learning with feed-forward control from an RNN-based surrogate model and it is implemented on an FPGA. The authors of [45] designed a fast feedback system based on reinforcement learning to suppress micro-bunch instabilities in the electron beam at KIT. Reference [46] provides an overview of how Bayesian optimizers can be applied online for efficient parameter tuning during real-time beam control, and the authors of [47] report a successful beam-tuning experiment at KEK using Bayesian optimisation.

Current and future challenges

Learning dynamics from data is in general straightforward compared to simulating the dynamics, which often require solving partial differential equations (PDEs) with complicated boundary conditions. Surrogate modelling is however a challenge for real-time accelerator applications due to highly nonlinear and fast dynamics, often in the kHz ranges and above.

In general, accelerator controls require fast hardware, such as FPGAs and time resolution below one microsecond. Since conventional PLC-based control systems lack the flexibility, scalability, and temporal precision needed, accelerator facilities develop their control systems in-house or rely on open-source controls communities. These communities have the potential to develop control system to meet the needs imposed by a higher level of automation. However, the engagement is limited. For example, the largest controls community, EPICS [48], currently has fewer than ten core developers. As a result, existing applications for storing and retrieving control system data offer limited functionality, and tools for life-cycle management of ML models within the control system are absent.

It is well known that data quality, i.e., how well data serves its intended purpose, is crucial for ML. In accelerator controls, maintaining high data quality is particularly challenging since accelerators are both fast and highly complex, with millions of process variables. The situation is worsened by a tendency to oversample slower process variables [49].

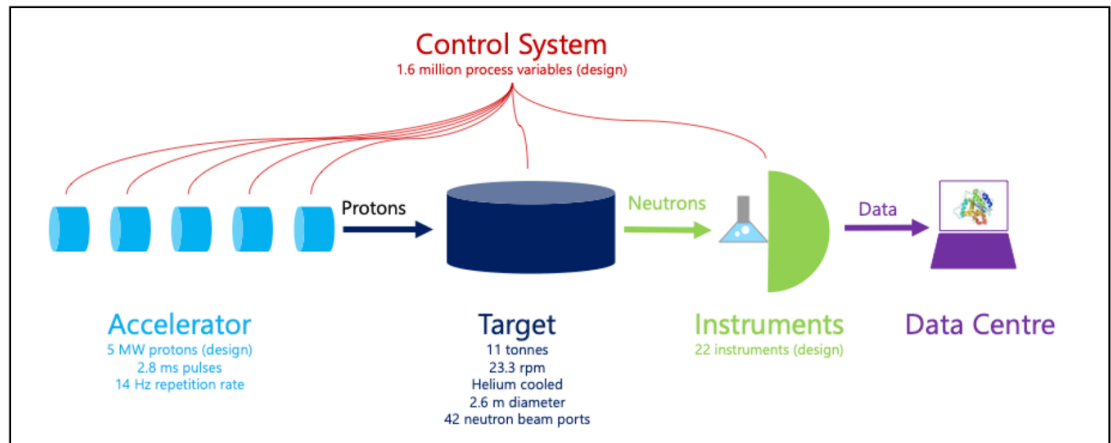


Figure 4. Schematic layout of the ESS machine with design parameters. The accelerator and target are used to create neutron beams, which are used for neutron scattering experiments in a variety of applied science fields.



Figure 5. Control racks at ESS. Some racks are empty since the high energy end of the accelerator, from where the picture is taken, has not yet been installed. The image gives an impression of the size of control systems at accelerator-based research facilities.

Despite the abundance of data within control system networks, the parameter space can remain sparsely populated due to the high dimensionality. Sparsity can also arise when certain process variables change only infrequently and require longer periods to adequately fill the parameter space. The coexistence of slow and fast processes therefore requires datasets that span long durations while still maintaining a high sampling rate.

High-power accelerators have high damage potential and require failsafe and fast $O(\mu\text{s})$ protection systems that can intervene within microseconds. Deploying fast ML models in this environment introduces additional risks, for example that a trained real-time ML model might exploit imperfections in the fail-safe protection system and drive the machine in a critical state.

Finally, ML for accelerator controls is foremost a research field in automatic controls and computer science applied to accelerators. However, external collaboration is time-consuming and requires efforts to document and transfer data from protected control systems, which involves both practical and organisational issues.

Advances in science and technology to meet challenges

Addressing the data quality challenge begins with establishing data-governance practises [49] and adapting a data-centric approach [50] to systematically and continuously improve data quality.

The issue of sparse data is to some extent also a data-quality problem, driven by the large number of process variables. Currently the number of process variables at ESS exceeds 10 million, which is considerably higher than the design value of 1.6 million. This can be partially mitigated by pruning the controls network and constraining operational ranges.

To address challenges in MLOps [51], the focus needs to be slightly shifted from ML modelling to the full workflow to develop, deploy, monitor and maintain ML models. Several initiatives have been started to develop platforms for MLOps, for example at CERN [52].

Accelerator-based research facilities have challenges and data that researchers and students are interested in, and therefore there are strong motivations for collaboration. To make data easily accessible for external parties, accelerator-based research facilities should share control system data [53]. For example, the dataset from Fermilab which was used in [44] is publicly available and well documented in [54]. Likewise, in [55] a publicly available dataset from ESS is described.

Finally, trained ML models will not be deployed without safety analysis and they will only be accepted for non-safety critical parts of the facility until they have proven safer than the operators, for example to detect potential faults. However, providing operators with a better understanding of the machine, with improved alarms, graphical representation of dependencies, virtual diagnostics and other augmentation techniques based on ML can help them in their daily work to operate the machine, if this information is accurate.

Concluding remarks

ML and other advanced methods will likely change the way accelerators are operated. Indeed, there are mathematical, structural, organisational, educational and safety-critical challenges to take on along the road towards autonomous accelerators. However, Accelerator-based research facilities offer unique environments to study and develop ML for process controls.

5. Fast ML for laser-plasma accelerators

Matthew Streeter and Charlotte Palmer

School of Mathematics and Physics, Queen's University Belfast, Belfast, United Kingdom

E-mail: m.streeter@qub.ac.uk and c.palmer@qub.ac.uk

Status

Laser plasma accelerators (LPAs) use ultra-short laser pulses focused to a high intensity ($\gtrsim 10^{18}$ W cm⁻²) to accelerate charged particles within a plasma. This field has developed rapidly over the last decade, demonstrating the capability to accelerate electron and ion beams to high energies over short (micron to centimetre) distances. A key metric of accelerators is the maximum achievable energy within a given distance. For laser-driven ion acceleration, the state of the art is acceleration of ~ 100 MeV protons using a high energy 200 J interacting with a ~ 100 nm thick plastic foils [56]. For laser-driven electron acceleration, electron energies up to 8 GeV [57] have been reached using a 20 cm long low-density plasma channel and a 31 J laser pulse through so called laser wakefield acceleration (LWFA).

Several proof-of-concept experiments have demonstrated the utility of these sources for near term applications. For example, LPA proton beams have been used for studies of radiobiology at ultra-high dose-rates [58], and for materials science [59]. X-rays generated in LWFAs by the accelerating electron beam have been used for radiography [60], tomography [61] and x-ray absorption spectroscopy [62]. Recent experiments have also demonstrated the potential of LWFA electron beams for use as compact drivers for x-ray FELs [63–65]. Owing to the multi-modal and highly tuneable nature of LPAs, there are many more applications which are being pursued at laser labs around the world [66].

With a longer-term view, plasma accelerators are an exciting prospect for ‘compact’ future high-energy particle colliders. This is due to the extremely high accelerating gradients that can be sustained in a plasma, making acceleration lengths many orders of magnitude shorter than ‘conventional’ radio-frequency accelerators. For example, GeV electrons have been accelerated in centimetre scale LPAs, in comparison to the 100 m scale conventional machines that achieve equivalent electron energies. In the long term, incorporating LPAs may offer a route to future colliders with a significantly reduced overall size [67], which is currently a major obstacle for extending the high energy frontier.

The dominant laser technology currently used in LPAs have repetition rates of 1–10 Hz at ~ 10 J per pulse, with a wall plug efficiency of $\lesssim 0.1\%$. Development of alternative laser technology promises to achieve this energy level at >1 kHz repetition rates and $>20\%$ efficiency [68] and is expected to dramatically improve the attractiveness of LPAs for many applications. This step change in repetition and data rates will present new opportunities and challenges for research in the field.

Current and future challenges

The non-linear physics at the core of extreme plasma accelerators presents challenges for control systems. The accelerator properties are highly sensitive to initial conditions and so development must focus on the stabilisation of laser drivers and plasma conditions. A demonstration experiment, using a specially stabilised laser system, has shown the potential to run an LWFA for >24 h continuously with electron beam energy stability on the few percent level [69]. The large datasets generated by modern multi-Hz facilities have also been used to reveal the sensitivity of electron beam parameters to the natural variation of the laser system [69, 70]. As the field matures, laser facilities are adopting automation into their operations and using real-time Bayesian Optimisation to efficiently optimise electron [71, 72], x-ray [71] and proton beams [73] from LPAs. These techniques have been mostly demonstrated on experimental systems operating at a few-Hz, due to the limitations of the laser systems. However, it will be necessary to increase the speed of the automation and feedback systems for optimal control of kHz plasma accelerators. The requirement for system stability and reliability has been recently highlighted by a project to develop a LPA injector for PETRA IV at DESY [74].

The challenging quality requirements on the outputs of LPAs motivates the need for active stabilisation of control parameters as well as the use of optimisation algorithms to tune the machines and identify stable regimes of operation [75]. DL techniques have been used to construct statistical models of LPAs [70, 76]. Even on few-Hz systems, such models have shown promise for predictive control systems which react to measurements taken at ~ 1 kHz to adjust for environmental conditions and correct the pointing of the laser system. Extending this to additional laser parameters (e.g. laser energy, wavefront, spectral phase) and for higher repetition rate laser systems will require fast diagnostic analysis and model inference on the sub-millisecond timescale. Future models may also be required to adapt for

drift in the machine states and correlations between controls to make long term operation more stable, robust, and reproducible. For instance, diagnosing spatiotemporal couplings (of critical importance in laser-plasma systems [77]) is complex and currently requires combinations of measurements and the use of ML models for inference [78]. This was reported to take on the order of 0.1 ms with a normal computer. Accelerated ML would provide much faster inference and enable such measurements at $\gtrsim 1$ kHz.

Operation of high power laser experiments is typically achieved through the independent control of multiple sub-systems with varying degrees of automation [79]. The next generation of high energy $\gtrsim 100$ Hz laser facilities will need to be more autonomous. High power laser systems are normally operated near their damage thresholds, which is achieved by careful balance of non-linear optical phenomena. It is relatively easy for a laser to enter states which cause significant damage to the system resulting in expensive repairs and lost beamtime. Therefore, it is paramount to incorporate reliable and robust control systems which can react fast enough to minimise this risk.

Advances in science and technology to meet challenges

Augmenting expert human operations with autonomous systems is likely required to ensure safe operations of future kHz high power laser systems. Indicators of problems in laser operation can manifest over just a few shots, which at high repetition rate would be far beyond the reaction time of a human operator. These autonomous systems will need to be reliable and obey both known and unknown constraints to prevent the overall machine, or any of its components, entering dangerous states. Such systems may also benefit from the capability to make predictions of damage before it occurs, monitoring for anomalies and responding before irreversible damage occurs, thereby minimising costly replacements and system downtime.

It will be necessary to adopt networked control systems, as is commonplace in large scale accelerator facilities, to fully automate experiments. To achieve this, the heterogeneity of existing systems will need to be overcome. Both the hardware and software systems will need to be capable of dealing with large data rates, performing data reduction and analysis, and communicating with remote high-performance computing (HPC) for physics modelling. Automation will require fast device control to minimise disruption to the machine outputs while performing optimisation and stabilisation tasks. Developments of radiation and electromagnetic pulse resistant components will be required, to ensure that systems are robust enough to survive the hostile environments created by high-intensity laser-plasma interactions.

Finally, LPAs are generally not fully diagnosed due to the difficulty of accessing information about the ultrafast microscopic accelerator dynamics as well as the destructive processes that occur when making use of the LPA outputs (e.g. loss of information of the incident electron spectrum in strong-field QED experiments [76]). ML modelling of such systems can suffer from aleatoric and epistemic uncertainties due to the random variation and drift of unobserved parameters. Accurately incorporating this into ML models will be required to maximise their utility. Approaches such as ‘CL’ may offer a way of maintaining optimal control of LPAs under such conditions, by endlessly adapting to changes to the machine state, e.g. degradation of optical components or thermal drifts in the laser system.

Concluding remarks

Laser plasma acceleration is an emerging technology which is moving into the realm of high-repetition rate operation (kHz) and high data-rates (Tb s^{-1}). Many techniques in automation can be adapted from developments at large scale accelerator facilities, but there are some unique characteristics that will require bespoke treatment. In particular, the highly non-linear physics of LPAs and of the laser systems themselves present challenges for autonomous control, especially in the presence of significant uncertainty about the machine states. Furthermore, the variety in the outputs of LPAs and within the end-user applications means that automation systems will have to incorporate configuration changes as efficiently as possible and in a user-friendly manner. The benefits of increasing the automation in this field, and incorporating ML for predictive and adaptive control, could be transformative in the utility of this technology by improving efficiency, stability, output quality and machine safety.

Acknowledgments

MJVS acknowledges support from the Royal Society URF-R1221874.

6. Fast ML for fusion simulation, optimisation, and control

Jonathan Citrin

Google DeepMind, London, United Kingdom

E-mail: citrin@google.com

Status

Fusion energy promises high-energy-density electricity and heat production with essentially limitless fuel reserves, inherent nuclear safety, and negligible impact on the environment [80]. The key societal importance of clean abundant energy, coupled with recent progress in the field, is currently fuelling intense interest and fusion investment in both the public and private spheres

Multiple approaches exist to achieve fusion energy. Laser fusion, a sub-branch of inertial confinement fusion (short fusion pulses at high plasma density), has recently achieved the historic milestone of net fusion gain, generating more fusion energy out than energy deposited into the fuel, in individual pulses at the National Ignition Facility in the US [81]. In magnetic confinement fusion, fusion occurs in lower density plasmas over long (ideally continuous) pulses. Various 2D and 3D magnetic configurations are studied, including stellarators, field-reversed-configurations, and tokamaks. Tokamaks are the most mature technology of this class. The next generation of tokamak experiments aims for net fusion power gain for the first time [82, 83]. While the challenges and simulation technologies outlined here share commonalities across the range of fusion approaches, the remainder of this review focuses on tokamaks.

Previous generations of tokamak device and experimental design have primarily relied on empirical scalings of plasma behaviour for determining device and plasma scenario configuration. However, continuous improvements in tokamak plasma simulation capabilities are bridging discrepancies between simulations and measurements, known as the *sim2real* gap. For example, as shown in figure 6, recent record-breaking fusion performance in the JET tokamak DT campaigns was successfully predicted by integrated modelling beforehand [84]. These advances increase confidence in theory-based extrapolations to performance in next-step devices, and enable a plethora of new simulation-based applications. In general, tokamak simulation has multiple use-cases. Interpretive simulations obtain information unavailable from measurements or help constraint measurements, for example calculating radial distributions of plasma heating and fuelling, and currents. Accurate simulations can provide physics understanding, validating theory-based models and untangling cause-and-effect in experiments through complex multi-physics simulations [85, 86]. The costs and risks of tokamak operation can be reduced by using simulation for experimental preparation (including inter-shot), performance optimisation including new scenario discovery, model-based controller design, and reactor design itself. In pulse-planning for next generation devices, pre-simulation is considered a requirement. ML can play a key role in accelerating simulation while maintaining accuracy, enabling these complex use-cases.

Current and future challenges

Tokamak integrated modelling is inherently multiscale and multiphysics, with multiple orders of magnitude in spatiotemporal scales between relevant physics processes. These include: magnetic equilibrium; transport arising from plasma collisions, magnetohydrodynamics and turbulence; heating and fuelling; line radiation arising from atomic and molecular interactions; and plasma material interaction. While high-fidelity direct numerical simulation is possible, this requires coupling multiple high-fidelity models at extreme (exascale) expense. Such workflows can provide ultimate ground truth for numerical verification of reduced models, but are not pragmatic for most use-cases.

A central challenge in tokamak simulation is thus bridging the conflicting constraints of model fidelity and tractability. Tokamak integrated modelling consists of coupled sets of transport PDEs, with reduced physics models calculating the various PDE coefficients [87]. The fidelity level of integrated modelling depends on the quality of these reduced models, with a trade-off between accuracy and tractability. Accurate but fast reduced physics models for these coefficients (related to plasma turbulence, heating, etc) are required. Multiple spatial domains must be simulated and coupled, comprising the plasma core (nested magnetic flux surfaces) which is amenable to a 1D radial approximation, the plasma edge region intersecting with the containment wall and requires a 2D description, and models for the plasma-facing-materials. In practice, workflows consisting of a single domain at a time are commonly used, although self-consistent coupled simulations are desirable. In general, this approach is suitable for experimental interpretation, gaining physics understanding, and extrapolation to future scenarios.

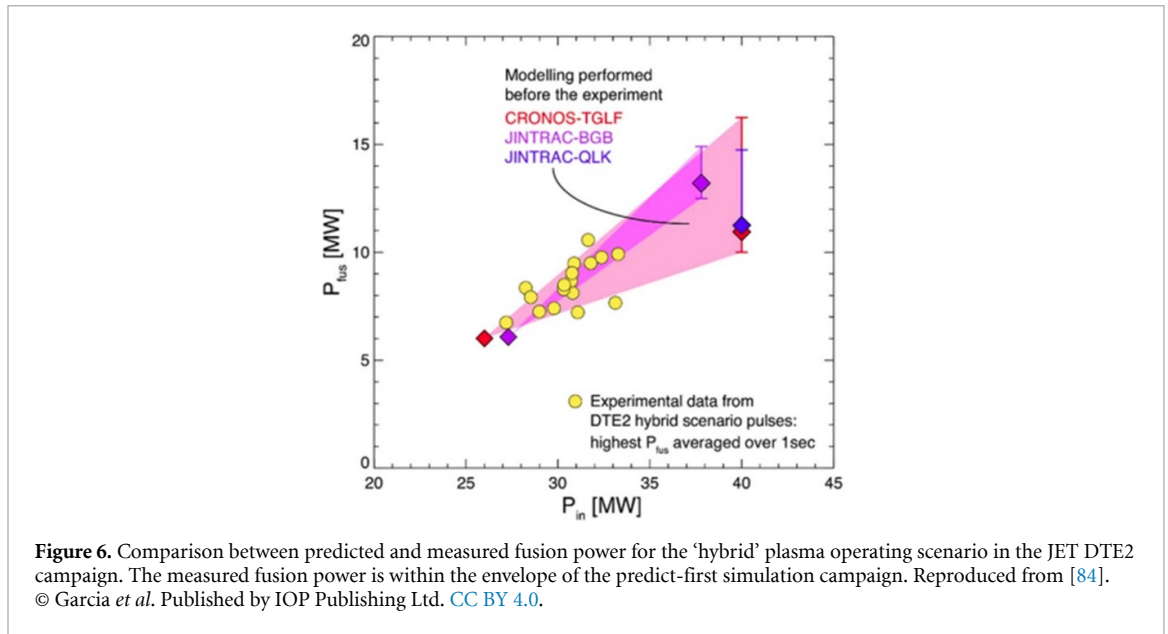


Figure 6. Comparison between predicted and measured fusion power for the ‘hybrid’ plasma operating scenario in the JET DTE2 campaign. The measured fusion power is within the envelope of the predict-first simulation campaign. Reproduced from [84]. © García *et al.* Published by IOP Publishing Ltd. CC BY 4.0.

However, accurate simulations of this type are typically too slow for many-query use cases such as pulse-design, optimisation, and controller-design.

Additional challenges arise within the context of realtime control. Fast and accurate simulation is required for theory-based nonlinear models in classical optimal control techniques such as model predictive control. Fast and accurate simulation is also required for realistic environments for agents in reinforcement learning approaches. Accurate state estimation for controllers would improve with higher availability of realtime diagnostics and low-latency event detection.

Advances in science and technology to meet challenges

A key method for improving tokamak multiphysics modelling speed is incorporating learned ML surrogates of integrated modelling physics components. These act as drop-in replacements of physics models incorporated in the transport PDEs. Furthermore, surrogates, e.g. based on neural networks, are differentiable, enabling their use within differentiable simulation for gradient-driven scenario optimisation and parameter identification. Surrogates are made using supervised learning techniques, learned from simulation databases generated in relevant parameter space using large-volume compute. General function approximators such as neural networks have the added advantage of being analytically differentiable functions, allowing incorporation into differentiable simulators for nonlinear PDE solvers and gradient-driven optimisation use-cases [88]. Prior knowledge of the input-output mapping characteristics of the physics models can be incorporated into the loss functions used for model training, improving evaluation accuracy. Numerous ML-surrogates for tokamak simulation have recently emerged. These include surrogates for neutral beam injection and current drive models [89], edge-transport-barrier formation [90], reduced turbulence models [90–92], and equilibrium [93]. These models were developed with different tokamak parameter ranges of validity, and cannot necessarily be incorporated into a single simulation. A concerted community effort to combine datasets and efforts, enabling a wide selection of ML-surrogates to be used in a single simulation for a given device, is desirable.

Taking advantage of the fact that in surrogate model generation, the primary computational burden is relegated to the data generation phase, ML-surrogates can be developed of models at higher fidelity than those typically used in integrated modelling, leading to simulation both faster and more accurate than the state-of-the-art [94]. High-fidelity models themselves can also be accelerated by ML techniques, for example by on-the-fly surrogate generation in narrow parameter space in simulation workflows [95], developing surrogates of internal expensive physics [96], or accelerated solver methods [97]. These approaches would enable generation of higher quality datasets for surrogate generation, and/or enable more routine high-fidelity verification of reduced models.

For physics components where theory-based models cannot deliver a sufficiently small sim2real gap, or are too expensive to generate sufficiently informative datasets for surrogate model generation, a hybrid data and model-driven approach is necessary [98]. Wider access to tokamak data would significantly accelerate development, as often access to quality data is the bottleneck in successful ML applications. Rapid advances in generative AI may also significantly improve data-driven predictions, e.g. by

learning latent structures, diffusion methods, and sequence-to-sequence models such as large language models (LLMs) and structured state spacemodels.

For realtime control applications, fast and accurate simulation enabled by ML surrogates can form environments used for training control policies with reinforcement learning, for example extending a recent approach pioneered for tokamak magnetic control, to multiphysics [99]. For state estimation, the range of available realtime diagnostics can be extended by using ML for accelerating inference, e.g. tomographic inversion, collisional-radiative calculations, or constrained equilibrium reconstruction. Reduced latency of these calculations would make previously privileged information available to controllers and supervisory systems, improving the efficacy of control policies.

Concluding remarks

ML can play a key role in the development and control of future fusion reactors. Learned surrogates can provide fast and accurate simulation, needed for pulse design and optimisation, as well as simulation environments for learning control policies with reinforcement learning. A combination of high-volume compute, adaptive sampling and workflow automation is required to develop the range of necessary surrogates. Open data practices at facilities is key for simulation validation, and incorporation of data-driven approaches for simulation components which have an insufficient theory-driven sim2real gap. Classical high-fidelity simulation is still required to refine reduced-order-models and constrain missing physics. ML can extend the range of realtime diagnostics available for state estimation, providing more information that can be used by control policies. Together, these approaches enable physics understanding, optimal pulse design, general control policies, all on the path towards the operational goals required for achieving fusion energy.

7. In-network ML: inference at the speed of data

Changgang Zheng and Noa Zilberman

Department of Engineering Science, University of Oxford, Oxford, United Kingdom

E-mail: changgang.zheng and noa.zilberman@eng.ox.ac.uk

Background and status

ML was widely adopted and extensively utilized by data-intensive applications such as stream processing and cybersecurity. However, the high demands of ML workloads pose challenges that local accelerators (e.g. GPUs, FPGAs) struggle to address. To cope with the scale and volume of the data to process, distributed ML is often used.

Networking and ML have a complex relationship. ML services rely on the network to carry data and connect accelerators, but often the exchange of data between accelerators is a bottleneck in distributed ML [100, 101]. At the same time, network devices can carry more data than an accelerator card can process, since a single network switch can process more than 100 Tbps, two orders of magnitude more than current accelerators. One way to bridge the mismatch between networking and ML is in-network ML. To explain it, a brief introduction to network programmability is provided.

High end network devices are designed to process as many packets as possible, as quickly as possible, exceeding 10 billion packets per second (Bpps) with sub-microsecond latency. Over the last decade, many network devices have become programmable [102], combining the processing pipeline (referred to as data plane) with reconfigurable match tables. The most popular general architecture is called PISA, Protocol Independent Switch Architecture, shown in figure 7(c). The domain specific language commonly used is P4.

As shown in figure 7, the PISA architecture has three components: (i) a programmable parser for extracting header information from incoming packets, (ii) a match-action pipeline that looks up keys (e.g. headers) in tables and matches them with actions (e.g. output port), and (iii) a deparser for reconstructing packets. Current programmable network devices range from software switches to switch-ASICs, FPGAs, SmartNICs, and DPUs.

The programmability of network devices enabled in-network computing, the offloading of applications normally running on a host to run within network devices. Examples of In-network computing applications include telemetry, caching [103] and gradient aggregation [101]. It demonstrated order of magnitude improvement in throughput, latency, and power efficiency [104]. Accelerating ML using in-network computing is one promising direction. This includes [105] both Network-Assisted ML, meaning using the network to accelerate workloads running on traditional targets, and in-network ML.

In-network ML offloads the entire ML inference process to programmable network devices [105]. It trains the ML model on servers or standard accelerators, and then maps the trained model to the device's data plane for inference. While resource utilisation varies, programmable switches often use only half of their resources for basic network functions [106], enabling in-network ML to leverage remaining resources. Models scale up to 20 trees and 55 features, achieving high classification accuracy [107]. The three main benefits of in-network ML are location, latency and load (3Ls) [107]. Network devices are already part of the infrastructure that data goes through (location), they can infer data closer to its source (latency) and inference performance scales with network bandwidth, while reducing load on end-hosts (load). In-network ML can be deployed either on network devices that are already on route from the data source, or as a network-attached accelerator. The first requires no additional cost or latency, while the second benefits from more available processing resources.

Meeting the exascale challenges

In-network ML inference process has two major components: feature extraction, and mapping of a trained ML algorithm to the data plane.

Feature Extraction. The extraction of features from incoming data can be done on various levels of granularity. Packet-level features are extracted directly from incoming packets, and are stateless, meaning the data from one packet does not affect later packets. Flow-level features consider the data flowing between a source-destination pair, and are stateful, meaning that information is accumulated over time, such as the traffic volume of the flow. Aggregation level features are also stateful, but combine information based on different metrics, e.g., the statistics of combining multiple ports. File-level feature extraction is possible but limited to certain types of files (e.g., text) and may have reduced performance, for example due to recirculation of packets within the pipeline. While packet-level features can be extracted in the parser,

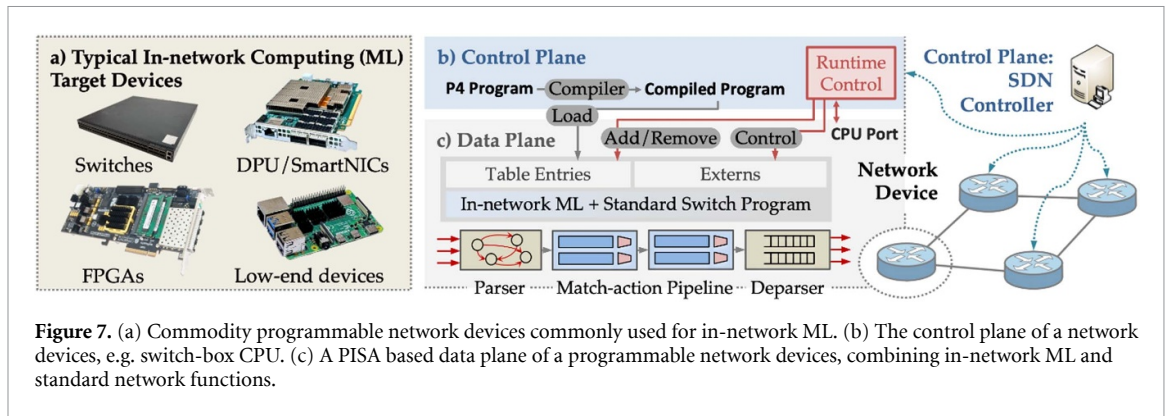


Figure 7. (a) Commodity programmable network devices commonly used for in-network ML. (b) The control plane of a network devices, e.g. switch-box CPU. (c) A PISA based data plane of a programmable network devices, combining in-network ML and standard network functions.

more complex feature extraction and feature engineering requires match-action pipeline stages and memory resources. One of the challenges for feature extraction is that packets arrival may be unordered, and information may travel through multiple routes.

ML Model Mapping. To maximise packet processing rate, network devices are constructed with limitations on programmability. For example, mathematical operations such as multiplication, and data types such as floating-point, are unavailable. Moreover, these are billions-of-transistors devices, therefore, to make their manufacturing feasible, user-available resources are limited in comparison with CPUs. For example, memory on a switch-ASIC is in the range of 100s of megabits, and the number of processing stages is e.g. 12 or 20. There is no native support for loops, and the closest equivalent, packet recirculation, has throughput implications. While some works have addressed these challenges through changes to the hardware [108] or by focusing on NICs [109], other successfully mapped inference models to off-the-shelf devices without loss of performance [105]. The mappings of ML models to the data plane can be divided into three groups [110]: Direct Mapping, where every step of the inference is translated into a stage in the pipeline, Lookup-based, where lookup tables are used for complex calculations, and Encode-based, where the feature space is sliced and encoded.

In-network ML Implementation. Implementing mapped ML models on target devices, and porting between different devices, requires significant effort. To solve this problem, Planter [110] offers rapid prototyping of in-network ML. Planter supports a range of ML models, architectures and target devices. It is modular, so users can extend the functionality. The framework also supports automatic parameter tuning, and several example use cases. Targets range from low-cost Raspberry Pi, through Alveo FPGA, to high performance Intel Tofino switch-ASIC. The demonstrated performance using ensemble tree models was 4.8 billion decisions per second.

Advances in science and technology to meet challenges

While in-network ML achieves high system performance, the size and complexity of the models that can be deployed are limited by resource constraints. This problem was identified by Microsoft as a key challenge for in-network computing [111], and Microsoft developed a resource elasticity solution [112] to address it.

As often only small or medium size models can be deployed on a network device [110, 113–117], ML performance can be lower than an unlimited size model running in a data centre. One solution is to distribute the model across multiple network devices, as demonstrated by DINC [106]. It has shown that ML models can be distributed across a data centre network or a wide area network, while coexisting with network functionality. A different solution, proposed in IIsy [107], is a hybrid deployment of in-network ML: deploying a small model on a network device, and a large model at the backend. This way, high confidence inference decisions are taken within the network, while only a fraction of the transactions needs to be inferred in a traditional manner. This enables both high ML performance and high system performance, while saving compute resources.

Data shift and model updates further challenge production deployments. P4Pir [118] suggests solving this problem by introducing a continuous update solution, focused on in-network ML for IoT.

Most in-network ML research focused so far on network traffic classification and cybersecurity [105, 107–109, 113, 115]. Examples that go beyond these use cases include load-balancing [119], predicting future stock prices [120], and image classification [121]. Still, the greatest challenge for in-network ML is wide adoption. This calls for extending its range of use cases, with data intensive and time-sensitive

applications being the immediate candidates. Real time processing of sensor data, for example, contextual streams from the GOTO telescope [122], from particle detectors at CERN's LHC [4], and from modern multi-Hz LWFA facilities [70], holds a significant potential.

Concluding remarks

In-network ML is an emerging technology, providing notable system performance benefits for ML inference. Available solutions enable deployment on commercial devices, supporting a range of ML models. Still, resource constraints remain a challenge, with hybrid and distributed deployments suggested as a solution.

There are opportunities for breakthrough research in the applications of in-network ML, benefiting the scientific community across disciplines, and now is the time to seize it.

Acknowledgments

This work was partly funded by VMware and Innovate UK (Project 10056403) as part of the SmartEdge EU project (Grant Agreement No. 101092908). We acknowledge support from Intel and NVIDIA. For the purpose of Open Access, the author has applied a CC BY public copyright license to any Author Accepted Manuscript version arising from this submission.

8. Accelerating traditional HPC using AI: a selective overview

Alexander Titterton

Graphcore, 11-19 Wine Street, Bristol BS1 2PH, United Kingdom

E-mail: alexandert@graphcore.ai

Status

Historically, HPC applications in fields such as high-energy particle physics, medical research and quantum chemistry have relied upon performing extensive calculations. We present a brief overview of some of the new applications in the intersection between AI and HPC, along with some of the challenges these aim to address.

Whilst modern supercomputing clusters typically feature a large number of processing cores, allowing for such applications to be highly parallelised, complex workloads can still take weeks or even months to complete. Furthermore, such applications often consist of a large codebase comprising thousands of lines of C++ code, in turn introducing the non-trivial challenge of code optimisation.

These challenges and the trade-off between simulation speed and accuracy motivate a new approach; a promising candidate for such being the introduction of AI/ML models into traditional HPC workflows. Experimentation with surrogate models—ML models which replace a numerically and computationally intensive part of an HPC workload—is becoming commonplace in various fields such as weather forecasting and high-energy particle physics.

Since an AI/ML model is simply learning to emulate the behaviour of a system, by producing an output which aims to better satisfy certain conditions in order to minimise the training loss, it need not be as computationally intensive as a simulation or analysis program written by hand. Furthermore, the mathematical operations which constitute these ML models are often well-suited to run on dedicated accelerator hardware in a massively-parallel fashion.

Mainstream Python-based DL frameworks such as PyTorch and TensorFlow support ML accelerators such as IPUs and GPUs natively. This has two key benefits: firstly, an application can easily and efficiently target dedicated hardware, offering huge speedups as well as better power efficiency (for example performance per Watt). Secondly, the amount of low-level code development required in order to produce a performance-optimised application is often dramatically reduced compared with more traditional HPC applications, thus allowing for more portable and reusable codebase which can more easily be shared and collaborated upon.

Using ML models in scientific HPC workloads is not only beneficial in terms of computational efficiency, however. There are many physical processes for which observation is possible, but we lack sufficient understanding of the underlying mechanisms to develop an accurate simulation. In such cases, ML models may be able to learn to emulate the behaviour of a system without requiring upfront an explicit description or set of rules which govern said system.

As shown in figure 8, AI-driven surrogate models may be used to reduce bottlenecks in HPC workloads by replacing a computationally expensive part of the process with an AI model, whilst preserving other elements of the pipeline. Other parts of the workload, for example data pre/post-processing and other simulation or analysis components are typically kept, rather than replacing the entire end-to-end process with an AI model, allowing for more fine-grained control over the outputs from each stage. In general, this surrogate model approach appropriately describes the implementation of each of the applications presented in this chapter.

Current and future challenges

As the scientific community advances, the ever-increasing complexity of data analysis, simulation and modelling drives the requirement for higher computing power. Since CPU clock speed has plateaued in recent years [123] parallel programming has become more popular, as has the use of dedicated hardware accelerators—though not without their own challenges and limitations. As was described in ‘The Hardware Lottery’ [124], often a proposed software solution is not that which best fits the problem one is trying to solve, but that which best suits the hardware one has on hand.

In the field of experimental high-energy particle physics, we see an enormous demand for computing resources corresponding to the sheer scale at which data is produced by LHC experiments. At the CMS experiment alone, the Phase-2 upgrade introduces a 4x increase in the recorded event size and a 7x increase in the event acceptance rate, leading to an overall increase of 20x in terms of the rate at which data is written to disk [125]. In order to ensure the event selection is efficient and robust, this in

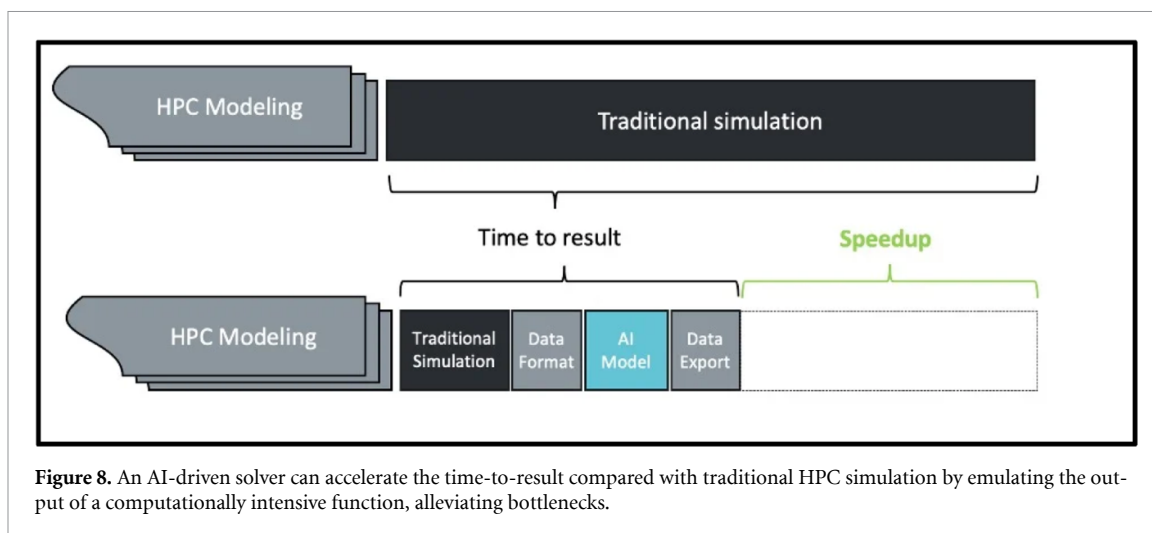


Figure 8. An AI-driven solver can accelerate the time-to-result compared with traditional HPC simulation by emulating the output of a computationally intensive function, alleviating bottlenecks.

turn requires faster and more complex ‘online’ event reconstruction and analysis tools. Such an increase in overall observed data also corresponds to a need to generate sufficient Monte Carlo simulated data, introducing further computing challenges.

The challenges in the field of medical drug discovery in the healthcare industry, however, are quite different. The average cost to develop a new drug lies in the region of \$2.6 Billion, often taking more than 10 years. Worse still, only a small percentage of new drugs are successful in moving from clinical trials to market [126]. Contrary to the case in high-energy particle physics, where accurate simulation of particle interactions is possible, the mapping of a protein’s structure to its function in medical treatment is not always well understood, and therefore an entirely new approach is required.

Nuclear fusion has potential to provide clean energy without consumption of fossil fuels. However only as recently as 2022 [127] has a fusion device been able to achieve a gain factor of greater than 1; i.e. produce more energy than was required to maintain the plasma steady state. Efficient power generation using a reactor such as a Tokamak requires a control system which is able to predict the evolution of the plasma state over time, which involves solving non-trivial PDEs. Solving such PDEs numerically is computationally intensive, yet must be performed in real time in order to be able to maintain a stable plasma.

Advances in science and technology to meet challenges

Whilst ML models can be more computationally efficient compared with HPC methods, achieving faster performance typically requires accelerator hardware. GPUs have long been repurposed for AI applications, but only recently has dedicated AI-first hardware begun to surface. An example is Graphcore’s IPU, which consists of a large number of independent processing cores along with on-chip memory [128], offering an alternative architecture to GPUs. Such alternatives can offer researchers new possibilities for more efficiently enhancing their workloads, whilst crucially supporting popular ML frameworks such as TensorFlow and PyTorch.

A variety of AI-driven approaches have been employed by the high-energy particle physics community. The use of generative models, for example GANs, enables accurate simulated data to be generated very quickly. A GAN comprises two ML models, with one trained to distinguish between generated and real data and the other trained to generate realistic data to fool the former. GANs are investigated for simulating di-jet production in [129], with the performance compared between CPUs, GPUs and IPUs. Recent work has shown further advances in the use of diffusion models for accelerated generation of simulated collision data [130]. Diffusion models are able to synthesize data by learning to iteratively reverse a fixed stochastic process that progressively corrupts a data sample into pure noise. Furthermore, following the release of ChatGPT, transformers have become popular in the AI community. The work in [131] introduces ParT (Particle Transformer) as a new benchmark for ML-driven jet tagging models.

In pharmaceutical research, *de novo* drug discovery is split into two main task categories: predictive (mapping protein structure to function) and generative (designing a protein to perform a desired function), with AI relevant in both cases. Transformer-based models are also proving popular here, with researchers starting with BERT [132], a model designed for English language, and replacing the vocabulary with one based upon the four bases of DNA sequences (A C G T). This approach has demonstrated

the ability to learn protein structural and functional properties [133]. However recent developments such as ProteinBERT [134] demonstrate how modifying models can better suit protein-related tasks.

An approach for solving PDEs is physics-informed neural networks (PINNs), whereby a PDE may be solved by training an ML model whose loss functions represent the PDE and its initial and boundary conditions. This type of application has potential to produce fast numerical solvers for PDEs which cannot be solved analytically, such as those which govern the plasma instability in a tokamak. PINNs are being investigated for such applications in nuclear fusion, however optimisation presents a challenge. The loss landscape for a PINN is often tumultuous, often resulting in the model getting stuck at a local optimum. To overcome such challenges, researchers in [135] investigate a hybrid methodology of training a PINN using labelled data and fine-tuning using an unsupervised approach.

Concluding remarks

As has been seen in this brief overview, the use of surrogate ML models to augment HPC applications in scientific research is an exciting and fast-progressing field. The ever-increasing complexity of computational challenges is clearly out-pacing the computational resources available, suggesting that in many cases in the near future traditional non-ML workloads simply may not suffice.

There remain numerous challenges, such as the explicability of a ‘black box’ neural network, and the need to build up relevant technical skills in the scientific community, particularly for students and junior researchers. In light of such challenges, close collaboration between academia and industry is more important than ever.

9. Inference speed is key to unleashing the potential of LLMs

Tobias Becker

Maxeler Technologies, a Groq Company, 3 Hammersmith Grove, London W6 0ND, United Kingdom

E-mail: tbecker@maxeler.com

Status

LLMs are a class of AI models that are trained on a vast amount of text data to generate language outputs that are coherent and natural-sounding. These models have become increasingly popular in recent years due to their ability to generate text that is often indistinguishable from human-written text. LLM models represent a significant breakthrough in natural language processing and have the potential to revolutionise many areas by providing a simple human language interface to computer systems that hold large amounts of data. Potential use cases include language translation, text summarisation, content generation, coding assistant, question answering, sentiment analysis, and many more.

LLMs are trained on massive datasets of text data, often from various sources such as books, articles, and websites. During training, the model learns to predict the next word in a sequence of text given the previous words. This process allows the model to learn the patterns and structures of language, such as grammar, syntax, and semantics.

The technology behind LLMs is based on a type of artificial neural network called a transformer [136], one of the most important innovations in the field of AI in recent years. Central to the transformer architecture is a technique called attention that allows the model to focus on different parts of the input text when generating each output word. This allows the model to capture long-range dependencies in the input text and generate coherent and natural-sounding output. Essentially, the attention mechanism offers a solution to the challenge of extreme combinatorial complexity when computing a longer sequence of words, therefore making it possible for the first time to generate high-quality language with computers. Despite this innovation, running inferences on an LLM is still computationally very demanding due to the very large size of the model which often has billions to hundreds of billions of parameters.

LLMs have gained widespread attention with the release of ChatGPT in November 2022 [137] and the field currently undergoes rapid innovation in many areas. Open-source models such as LLaMa have been released [138], and various techniques have been developed to customise LLMs without having to retrain the entire model from scratch [139]. While LLMs are currently limited in speed and accessibility, further advances in models and computer hardware are expected to make LLMs much more available to everyone.

Current and future challenges

There are several challenges associated with LLMs and many of those relate to the sheer size of the models themselves. Training LLMs is extremely compute intensive and requires large amounts of high-quality training data. GPT-4, a very large model with over one trillion parameters that was created by OpenAI, is estimated to have taken 90–100 d of training time on a cluster with 25 000 Nvidia A100 GPUs at an estimated cost of over \$100 M [139]. Training also requires access to a large number of data sources with high-quality text. Training LLMs is therefore an activity that is only accessible to a small number of large institutions, and further increasing model sizes will increase this challenge further. It is also unclear if ever increasing model sizes are necessary. Recent work suggests indeed that increasing the model size will lead to better performance [140] while other research disagrees and indicates that some large models will perform worse at certain tasks than smaller ones [141]. Another challenge lies in measuring the quality of an output given by an LLM to which there is currently no single agreed metric.

Another significant challenge with LLMs relates to the fact that these models inevitably have limited domain knowledge. Even if a large amount of training data is used, the model will have limited knowledge of specialised areas such as law or medicine. Furthermore, the information embedded in the model is static and does not capture any new information. LLMs can therefore be limited in their ability to produce usable outputs in specialised use cases or to generalize to new, unseen data such as giving answers related to recent events. Due to the high cost of training, retraining an entire LLM for a specialised use case or new data will not be feasible. Finally, the last challenge relates to the compute complexity of the LLM inference, and the speed needed for practical use cases. For example, LLaMa 270B can require, depending on implementation parameters, over 100GFlops to produce a single output token (roughly the equivalent of a word). At the same time, the computation cannot be heavily parallelised

because of the sequential dependencies between the tokens in the model. This means that real-world performance of current LLM systems is often sluggish, with single digit to tens of tokens per second being produced.

Advances in science and technology to meet challenges

The entire field of LLMs currently undergoes rapid development and innovation to address the various challenges mentioned above. One area of research addresses the complexity of the training process with the goal to reduce the overall compute time needed for training [142], but the overall effort for training remains very high. There are also several techniques to reduce the compute complexity during inference such as quantisation, knowledge distillation, or pruning. A large body of research addresses modifications and optimisations of the transformer architecture itself. For example, mixture of expert models [143], multi-head attention for decoding [144], or sparse transformer models [145] all address the compute complexity in the transformer model during the inference process.

There are several techniques that address the limitations of fixed, pre-trained LLMs. Fine-tuning is a training technique that enables adapting a pre-trained model to a specific task or specialised domain, e.g. language translation, question answering, medical knowledge. It is a much more efficient and lightweight process than training a full model from scratch and involves feeding a small amount of task-specific data to the model and adjusting some of the model's parameters. Low rank adaptation is an example of an efficient fine-tuning technique [146]. Other approaches incorporate additional data sources into the LLM. Retrieval-augmented generation extends a pre-trained LLM with retrieval model that can feed in relevant information from a separate database and augment the generation process. This is a way of adding new or specialised information that is not embedded in the pre-trained LLM [147].

In order the run LLMs are the speeds needed for practical applications, suitable compute hardware is needed. At present, most LLM services run on GPU hardware resulting in inference speeds of tens of tokens per seconds. However, due to the sequential properties of LLMs, it is difficult to optimise throughput by parallelising across many GPUs. As an alternative, more specialised hardware can be used. The AWS Inferentia2 inference accelerator can run various Llama 2 models cost efficiently but faces similar parallelisation challenges as GPUs [148]. Recently, much higher throughput of several hundreds of tokens per second have been reported when using GroqChip, an AI inference accelerator that can handle the sequential properties of LLMs much better than conventional CPUs or GPUs. GroqChip uses a large amount of on-chip SRAM which enables fast, low-latency inference. A dedicated chip-to-chip interconnect between the multiple GroqChips offers scalability for running large models with very high throughput. Results are rapidly evolving and point towards continued improvement on inference speeds across various vendors [149].

Concluding remarks

LLMs have the potential to provide many benefits across many different sectors and application domains by providing a human language interface to powerful computers that have access to large amounts of data. Adding a human language interface to computers could be similarly transformative to the introduction of graphical user interfaces many decades ago that made computers much more usable by a wider audience. However, current LLM technology faces many challenges and many relate to the sheer size of these models, most notably the high computational effort to train and run LLMs. Training is extremely costly and running inference on LLMs is also costly and too slow for many practical applications. At present, the field is rapidly evolving to find solutions to these challenges. New fine-tuning techniques reduce the need for full model training. Advances in models and techniques combined with new specialised AI accelerator hardware are improving inference speeds. Together, these developments will overcome current limitations and soon lead to LLM performance good enough for a wide range of real-world applications.

Acknowledgments

The Guest Editors of this Roadmap would like to thank the workshop scientific committee.

Fast Machine Learning Workshop 2023 Scientific Committee:

Thea Aarrestad (ETH Zurich)
Javier Duarte (UCSD)
Phil Harris (MIT)
Burt Holzman (Fermilab)

Scott Hauck (U. Washington)
Shih-Chieh Hsu (U. Washington)
Sergo Jindariani (Fermilab)
Mia Liu (Purdue University)
Allison McCarn Deiana (Southern Methodist University)
Mark Neubauer (U. Illinois Urbana-Champaign)
Jennifer Ngadiuba (Fermilab)
Maurizio Pierini (CERN)
Sioni Summers (CERN)
Alex Tapper (Imperial College London)
Nhan Tran (Fermilab)

Data availability statement

No new data were created or analysed in this study.

ORCID iDs

Sioni Summers  0000-0003-4244-2061
Alex Tapper  0000-0003-4543-864X
Thea Klæboe Årrestad  0000-0002-7671-243X
Karin Rathsmann  0009-0005-0715-8905
Matthew Streeter  0000-0001-9086-9831
Charlotte Palmer  0000-0002-3762-0380
Changgang Zheng  0000-0003-1894-722X
Noa Zilberman  0000-0002-3655-2873
Tobias Becker  0000-0001-5939-6757

References

- [1] Duarte J *et al* 2018 Fast inference of deep neural networks in FPGAs for particle physics *J. Instrum.* **13** P07027
- [2] Summers S *et al* 2020 Fast inference of boosted decision trees in FPGAs for particle physics *J. Instrum.* **15** P05026
- [3] A3D3 Institute 2024 AI to accelerate scientific discovery (A3D3) (available at: <https://a3d3.ai/about/>) (Accessed 10 January 2024)
- [4] CMS Collaboration 2020 The Phase-2 Upgrade of the CMS Level-1 Trigger, CERN, Geneva *Tech. Rep. CERN-LHCC-2020-004, CMS-TDR-021* (CERN) (available at: <https://cds.cern.ch/record/2714892>)
- [5] DeZoort G, Thais S, Duarte J, Razavimaleki V, Atkinson M, Ojalvo I, Neubauer M and Elmer P 2021 Charged particle tracking via edge-classifying interaction networks *Comput. Softw. Big Sci.* **5** 26
- [6] Elabd A *et al* 2022 Graph neural networks for charged particle tracking on FPGAs *Front. Big Data* **5** 828666
- [7] Brown C E, Bundock A, Komm M, Loncar V, Pierini M, Radburn-Smith B C, Shtipliyski A, Summers S P, Dancu J-S and Tapper A 2023 Neural network-based primary vertex reconstruction with FPGAs for the upgrade of the CMS Level-1 trigger system *Tech. Rep. CMS-CR-2022-018* (CERN) (available at: <https://cds.cern.ch/record/2801638>)
- [8] Iiyama Y *et al* 2021 Distance-weighted graph neural networks on FPGAs for real-time particle reconstruction in high energy physics *Front. Big Data* **3** 598927
- [9] Schaefer S, Brown C, Hoang D, Summers S and Wuchterl S 2025 Advancing the CMS Level-1 trigger: jet tagging with deepSets at the HL-LHC (arXiv:2509.24371)
- [10] Que Z, Loo M, Fan H, Pierini M, Tapper A and Luk W 2022 Optimizing graph neural networks for jet tagging in particle physics on FPGAs *Proc. 32nd Int. Conf. on Field-Programmable Logic and Applications* (<https://doi.org/10.1109/FPL57034.2022.00057>)
- [11] Wojcicki F, Que Z, Tapper A D and Luk W 2022 Accelerating transformer neural networks on FPGAs for high energy physics experiments *Proc. 2022 IEEE Int. Conf. on Field-Programmable Technology (ICFPT)* (<https://doi.org/10.1109/ICFPT56656.2022.9974463>)
- [12] Gandrakota A 2024 Real-time anomaly detection at the L1 trigger of CMS experiment (arXiv:2411.19506)
- [13] Cohen M M *et al* (ATLAS Collaboration) 2025 GELATO: a generic event-level anomalous trigger option for ATLAS—slides (available at: <https://cds.cern.ch/record/2938881>)
- [14] CMS Collaboration 2023 Electron reconstruction and identification in the CMS Phase-2 Level-1 trigger (available at: <https://cds.cern.ch/record/2868782>)
- [15] Reikher D 2025 Leveraging machine learning for enhanced FPGA- based tau triggering and combined $H \rightarrow (bb/cc)$ analysis in the ATLAS experiment (<https://doi.org/10.17181/93sta-ptr08>)
- [16] CMS Collaboration 2023 Continual learning in the CMS Phase-2 Level-1 trigger (available at: <https://cds.cern.ch/record/2859651>)
- [17] Kulkarni S R *et al* 2023 On-sensor data filtering using neuromorphic computing for high energy physics experiments (arXiv:2307.11242[cs.NE])
- [18] Dickinson J *et al* 2023 Smartpixels: towards on-sensor inference of charged particle track parameters and uncertainties (arXiv:2312.11676 [hep-ex])
- [19] Di Guglielmo G *et al* 2021 A reconfigurable neural network ASIC for detector front-end data compression at the HL-LHC *IEEE Trans. Nucl. Sci.* **68** 2179–86
- [20] Yue A, Jia H and Gonski J 2025 Variational autoencoders for at-source data reduction and anomaly detection in high energy particle detectors *Mach. Learn.: Sci. Technol.* **6** 035017

- [21] Rinck P A 2019 Magnetic resonance in medicine: a critical introduction *The Basic Textbook of the European Magnetic Resonance Forum* (BoD Book on Demand)
- [22] Leiner T, Rueckert D, Suinasiaputra A, Baeßler B, Nezafat R, İşgum I and Young A A 2019 Machine learning in cardiovascular magnetic resonance: basic concepts and applications *J. Cardiovasc. Magn. Reson.* **21** 1–14
- [23] Blansit K, Retson T, Masutani E, Bahrami N and Hsiao A 2019 Deep learning–based prescription of cardiac MRI planes *Radiology* **1** e180069
- [24] Qin C, Schlemper J, Caballero J, Price A N, Hajnal J V and Rueckert D 2018 Convolutional recurrent neural networks for dynamic MR image reconstruction *IEEE Trans. Med. Imaging* **38** 280–90
- [25] Qin C, Duan J, Hammernik K, Schlemper J, Küstner T, Botnar R, Prieto C, Price A N, Hajnal J V and Rueckert D 2021 Complementary time-frequency domain networks for dynamic parallel MR image reconstruction *Magn. Reson. Med.* **86** 3274–91
- [26] Bai W et al 2018 Automated cardiovascular magnetic resonance image analysis with fully convolutional networks *J. Cardiovasc. Magn. Reson.* **20** 1–12
- [27] Qin C, Bai W, Schlemper J, Petersen S E, Piechnik S K, Neubauer S and Rueckert D 2018 Joint learning of motion estimation and segmentation for cardiac MR image sequences *Proc., Part II. Medical Image Computing and Computer Assisted Intervention–MICCAI 2018: 21st Int. Conf. (Granada, Spain, 16–20 September 16–20 2018)* vol 11 (Springer) pp 472–80
- [28] Qin C, Shi B, Liao R, Mansi T, Rueckert D and Kamen A 2019 Unsupervised deformable registration for multi-modal images via disentangled representations *Int. Conf. on Information Processing in Medical Imaging* (Springer) pp 249–61
- [29] Topol E 2019 The topol review *Preparing the Healthcare Workforce to Deliver the Digital Future* pp 1–48
- [30] Glocker B, Robinson R, Castro D C, Dou Q and Konukoglu E 2019 Machine learning with multi-site imaging data: an empirical study on the impact of scanner effects *Medical Imaging Meets NeurIPS Workshop*
- [31] Littlejohns T J et al 2020 The UK Biobank imaging enhancement of 100,000 participants: rationale, data collection, management and future directions *Nat. Commun.* **11** 2624
- [32] Kondylakis H et al 2026 A review of methods for trustworthy AI in medical imaging: the FUTURE-AI Guidelines *IEEE J. Biomed. Health Inform.* **30** 2299–315
- [33] Antun V, Renna F, Poon C, Adcock B and Hansen A C 2020 On instabilities of deep learning in image reconstruction and the potential costs of AI *Proc. Natl Acad. Sci.* **117** 30088–95
- [34] Shurrab S and Duwairi R 2022 Self-supervised learning methods and applications in medical imaging analysis: a survey *PeerJ Comput. Sci.* **8** e1045
- [35] Kazerouni A, Aghdam E K, Heidari M, Azad R, Fayyaz M, Hacihaliloglu I and Merhof D 2023 Diffusion models in medical imaging: a comprehensive survey *Med. Image Anal.* **88** 102846
- [36] Zhang S and Metaxas D 2023 On the challenges and perspectives of foundation models for medical image analysis *Med. Image Anal.* **91** 102996
- [37] Kaissis G et al 2021 End-to-end privacy preserving deep learning on multi-institutional medical imaging *Nat. Mach. Intell.* **3** 473–84
- [38] Begoli E, Bhattacharya T and Kusnezov D 2019 The need for uncertainty quantification in machine-assisted medical decision making *Nat. Mach. Intell.* **1** 20–23
- [39] Singh A, Sengupta S and Lakshminarayanan V 2020 Explainable deep learning models in medical image analysis *J. Imaging* **6** 52
- [40] Vlontzos A, Rueckert D and Kainz B 2022 A review of causality for learning algorithms in medical image analysis *Mach. Learn. Biomed. Imaging* **1** 1–17
- [41] Gleyzer S, Carminati F, Wonju G and Perret-Gallix D 2018 The rise of deep learning (CERN Courier (available at: <https://cerncourier.com/a/the-rise-of-deep-learning/>)) (Accessed 12 January 2024)
- [42] PSI 2019 2nd ICFA mini-workshop on machine learning for charged particle accelerators (Indico) (available at: <https://indico.psi.ch/event/6698/>) (Accessed 1 2024)
- [43] 2022 *18th Int. Conf. on Acc. and Large Exp. Physics Control Systems* (JACoW Publishing)
- [44] St. John J et al 2021 Real-time artificial intelligence for accelerator control *Phys. Rev. Accel. Beams* **24** 104601
- [45] Boltz T, Brosi M, Bründermann E, Haerer B, Kaiser P, Pohl C, Schreiber P, Yan M, Asfour T and Müller A-S 2020 Feedback design for control of the micro-bunching instability based on reinforcement learning *CERN Yellow Reports: Conf. Proc.* (CERN)
- [46] Roussel R et al 2024 Bayesian optimization algorithms for accelerator physics *Phys. Rev. Accel. Beams* **27** 084801
- [47] Mitsuka G, Kato S, Iida N, Natsui T and Satoh M 2024 Machine-learning approach for operating electron beam at KEK electron/positron injector linac *Phys. Rev. Accel. Beams* **27** 084601
- [48] Experimental Physics and Industrial Control System (EPICS) Home page (available at: <https://epics-controls.org/>)
- [49] Källström C 2025 Data quality and quantity for machine learning at the European spallation source
- [50] Jarrahi M H, Memariani A and Guha S 2023 The principles of data-centric AI *Assoc. Comput. Mach.* **66** 84–92
- [51] Kreuzberger D, Kühl N and Hirschl S 2023 Machine learning operations (MLOps): overview, definition, and architecture *IEEE Access* **11** 31866–79
- [52] de Martel J-B, Gorbosov R and Madysa N 2021 Machine learning platform: deploying and managing models in the CERN control system *18th Int. Conf. on Acc. and Large Exp. Physics Control Systems (Shanghai)*
- [53] Runeson P, Olsson T and Linåker J 2021 Open data ecosystems—an empirical investigation into an emerging industry collaboration concept *J. Syst. Softw.* **182** 111088
- [54] Kafkes D and St. John J B 2021 A dataset for accelerator control systems data *Data* **6** 42
- [55] Mogensen S W, Rathsman K and Nilsson P 2024 Causal discovery in a complex industrial system: a time series benchmark *Proc. Third Conf. on Causal Learning and Reasoning (Los Angeles)*
- [56] Higginson A et al 2018 Near-100 MeV protons via a laser-driven transparency-enhanced hybrid acceleration scheme *Nat. Commun.* **9** 724
- [57] Gonsalves A J et al 2019 Petawatt laser guiding and electron beam acceleration to 8 GeV in a laser-heated capillary discharge waveguide *Phys. Rev. Lett.* **122** 084801
- [58] Kroll F et al 2022 Tumour irradiation in mice with a laser-accelerated proton beam *Nat. Phys.* **18** 316–22
- [59] Barberio M et al 2018 Laser-accelerated particle beams for stress testing of materials *Nat. Commun.* **9** 372
- [60] Hussein A E et al 2019 Laser-wakefield accelerators for high-resolution x-ray imaging of complex microstructures *Sci. Rep.* **9** 3249
- [61] Cole J M et al 2016 Tomography of human trabecular bone with a laser-wakefield driven x-ray source *Plasma Phys. Control. Fusion* **58** 014008

- [62] Kettle B et al 2019 Single-shot multi-keV x-ray absorption spectroscopy using an ultrashort laser wakefield accelerator source *Phys. Rev. Lett.* **123** 254801
- [63] Wang W et al 2021 Free-electron lasing at 27 nanometres based on a laser wakefield accelerator *Nature* **595** 516–20
- [64] Labat M et al 2023 Seeded free-electron laser driven by a compact laser plasma accelerator *Nat. Photon.* **17** 150–6
- [65] Barber S K et al 2025 Greater than 1000-fold gain in a free-electron laser driven by a laser-plasma accelerator with high reliability *Phys. Rev. Lett.* **135** 055001
- [66] Danson C N et al 2019 Petawatt and exawatt class lasers worldwide *High Power Laser Sci. Eng.* **7** e54
- [67] Foster B, D’Arcy R and Lindström C A 2023 A hybrid, asymmetric, linear Higgs factory based on plasma-wakefield and radio-frequency acceleration *New J. Phys.* **25** 093037
- [68] Kiani L et al 2023 High average power ultrafast laser technologies for driving future advanced accelerators *J. Instrum.* **18** T08006
- [69] Maier A R et al 2020 Decoding sources of energy variability in a laser-plasma accelerator *Phys. Rev. X* **10** 031039
- [70] Kirchen M et al 2021 Optimal beam loading in a laser-plasma accelerator *Phys. Rev. Lett.* **126** 174801
- [71] Shaloo R J et al 2020 Automation and control of laser wakefield accelerators using Bayesian optimization *Nat. Commun.* **11** 6355
- [72] Jalas S, Kirchen M, Messner P, Winkler P, Hübner L, Dirkwinkel J, Schnepf M, Lehe R and Maier A R 2021 Bayesian optimization of a laser-plasma accelerator *Phys. Rev. Lett.* **126** 104801
- [73] Loughran B et al 2023 Automated control and optimisation of laser driven ion acceleration *High Power Laser Sci. Eng.* **11** 1–11
- [74] Agapov I et al 2025 *The Plasma Injector for PETRA IV: Enabling Plasma Accelerators for Next-generation Light Sources. Conceptual Design Report* (Deutsches Elektronen-Synchrotron DESY)
- [75] Döpp A, Eberle C, Howard S, Irshad F, Lin J and Streeter M 2023 Data-driven science and machine learning methods in laser-plasma physics *High Power Laser Sci. Eng.* **11** e55
- [76] Streeter M J V et al 2023 Laser wakefield accelerator modelling with variational neural networks *High Power Laser Sci. Eng.* **11** e9
- [77] Jeandet A et al 2022 Survey of spatio-temporal couplings throughout high-power ultrashort lasers *Opt. Express* **30** 3262–88
- [78] Howard S, Esslinger J, Weiße N, Schröder J, Eberle C, Wang R H W, Karsch S, Norreys P and Döpp A 2025 Single-shot spatio-temporal vector field measurements of petawatt laser pulses *Nat. Photon.* **19** 898–905
- [79] Feister S et al 2023 Control systems and data management for high-power laser facilities *High Power Laser Sci. Eng.* **11** 1–31
- [80] Sadiq-Zada E R, Gatto A and Weißnicht Y 2023 Back to the future: revisiting the perspectives on nuclear fusion and juxtaposition to existing energy sources *Energy* **290** 129150
- [81] Zylstra A B et al 2022 Burning plasma achieved in inertial fusion *Nature* **601** 524–48
- [82] Bigot B 2022 Preparation for assembly and commissioning of ITER *Nucl. Fusion* **62** 042001
- [83] Creely A J et al 2020 Overview of the SPARC tokamak *J. Plasma Phys.* **86** 865860502
- [84] Garcia J et al 2023 Modelling performed for predictions of fusion power in JET DTE2: overview and lessons learnt *Nucl. Fusion* **63** 112003
- [85] Marin M et al 2021 Multiple-isotope pellet cycles captured by turbulent transport modelling in the JET tokamak *Nucl. Fusion* **61** 036042
- [86] Rodriguez-Fernandez P et al 2018 Explaining cold-pulse dynamics in tokamak plasmas using local turbulent transport models *Phys. Rev. Lett.* **120** 075001
- [87] Poli F M 2018 Integrated Tokamak modeling: when physics informs engineering and research planning *Phys. Plasmas* **25** 055602
- [88] Van Mulders S, Felici F, Sauter O, Citrin J, Ho A, Marin M and van de Plassche K L 2021 Rapid optimization of stationary tokamak plasmas in RAPTOR: demonstration for the ITER hybrid scenario with neural network surrogate transport model QLKNN *Nucl. Fusion* **61** 086019
- [89] Boyer M D, Kaye S and Erickson K 2019 Real-time capable modeling of neutral beam injection on NSTX-U using neural networks *Nucl. Fusion* **59** 056008
- [90] Meneghini O et al 2017 Self-consistent core-pedestal transport simulations with neural network accelerated models *Nucl. Fusion* **57** 086034
- [91] van de Plassche K L et al 2020 Fast modeling of turbulent transport in fusion plasmas using neural networks *Phys. Plasmas* **27** 022310
- [92] Ho A et al 2023 Predictive JET current ramp-up modelling using QuaLiKiz-neural-network *Nucl. Fusion* **63** 066014
- [93] Joung S, Ghim Y-C, Kim J, Kwak S, Kwon D, Sung C, Kim D, Kim H-S, Bak J G and Yoon S W 2023 GS-DeepNet: mastering tokamak plasma equilibria with deep neural networks and the Grad-Shafranov equation *Sci. Rep.* **13** 15799
- [94] Citrin J, Trochim P, Goerler T, Pfau D, van de Plassche K L and Jenko F 2023 Fast transport simulations with higher-fidelity surrogate models for ITER *Phys. Plasmas* **30** 062501
- [95] Rodriguez-Fernandez P, Howard N T and Candy J 2022 Nonlinear gyrokinetic predictions of SPARC burning plasma profiles enabled by surrogate modeling *Nucl. Fusion* **62** 076036
- [96] Miller S T, Roberts N V, Bond S D and Cyr E C 2022 Neural-network based collision operators for the Boltzmann equation *J. Comput. Phys.* **470** 111541
- [97] Tanyu D N, Ning J, Freudenberg T, Heilenkötter N, Rademacher A, Iben U and Maass P 2023 Deep learning methods for partial differential equations and related parameter identification problems *Inverse Problems* **39** 103001
- [98] Hatfield P W et al 2021 The data-driven future of high-energy-density physics *Nature* **593** 351–61
- [99] Degraeve J et al 2022 Magnetic control of tokamak plasmas through deep reinforcement learning *Nature* **602** 414–9
- [100] Luo L, West P, Nelson J, Krishnamurthy A and Ceze L 2020 Plink: discovering and exploiting locality for accelerated distributed training on the public cloud *Proc. Mach. Learn. Syst.* **2** 82–97
- [101] Sapio A, Canini M, Chen-Yu H, Nelson J, Kalnis P, Kim C, Krishnamurthy A, Moshref M, Ports D and Richtarik P 2021 Scaling distributed machine learning with In-Network aggregation *18th USENIX Symp. on Networked Systems Design and Implementation (NSDI 21)* pp 785–808
- [102] Bosshart P, Gibb G, Kim H-S, Varghese G, McKeown N, Izzard M, Mujica F and Horowitz M 2013 Forwarding metamorphosis: fast programmable match-action processing in hardware for SDN *ACM SIGCOMM Comput. Commun. Rev.* **43** 99–110
- [103] Jin X, Xiaozhou L, Zhang H, Soule R, Lee J, Foster N, Kim C and Stoica I 2017 NetCache: balancing key-value stores with fast in-network caching *Proc. 26th Symp. on Operating Systems Principles* pp 121–36
- [104] Tokusashi Y, Dang H T, Pedone F, Soule R and Zilberman N 2019 The case for in-network computing on demand *Proc. Fourteenth EuroSys Conf. 2019* pp 1–16
- [105] Zheng C, Hong X, Ding D, Vargaftik S, Ben-Itzhak Y and Zilberman N 2023 In-network machine learning using programmable network devices: a survey *IEEE Commun. Surv. Tutorials* **26** 1171–200

- [106] Zheng C, Tang H, Zang M, Hong X, Feng A, Tassiulas L and Zilberman N 2023 DINC: toward distributed in-network computing *Proc. ACM Netw.* **1** 1–25
- [107] Zheng C, Xiong Z, Bui T T, Kaupmees S, Bensoussane R, Bernabeu A, Vargaftik S, Ben-Itzhak Y and Zilberman N 2024 IIsy: hybrid in-network classification using programmable switches *IEEE/ACM Trans. on Networking* pp 2555–70
- [108] Swamy T, Rucker A, Shahbaz M, Gaur I and Olukotun K 2022 Taurus: a data plane architecture for per-packet ML *Proc. 27th ACM Int. Conf. on Architectural Support for Programming Languages and Operating Systems* pp 1099–114
- [109] Siracusano G, Galea S, Sanvito D, Malekzadeh M, Antichi G, Costa P, Haddadi H and Bifulco R 2022 Re-architecting traffic analysis with neural network interface cards *19th USENIX Symp. on Networked Systems Design and Implementation (NSDI 22)* pp 513–33
- [110] Zheng C, Zang M, Hong X, Perreault L, Bensoussane R, Vargaftik S, Ben-Itzhak Y and Zilberman N 2024 Planter: rapid prototyping of in-network machine learning inference *ACM SIGCOMM Comput. Commun. Rev.* **54** 2–21
- [111] Victor Bahl Unlocking the potential of in-network computing for telecommunication workloads (Website) (available at: <https://azure.microsoft.com/en-us/blog/unlocking-the-potential-of-in-network-computing-for-telecommunication-workloads/>) (Accessed November 2023)
- [112] Kim D, Liu Z, Zhu Y, Kim C, Lee J, Sekar V and Seshan S 2020 Tea: enabling state-intensive network functions on programmable switches *Proc. 2020 ACM SIGCOMM Conf.* pp 90–106
- [113] Akem A T-J, Gucciardo M and Fiore M 2023 Flowrest: practical flow-level inference in programmable switches with random forests *IEEE INFOCOM 2023-IEEE Conf. on Computer Communications*
- [114] Coralie B-G, Meier R, Dietmüller A, Bühler T and Vanbever L 2019 pforest: in-network inference with random forests (arXiv:1909.05680)
- [115] Jong-Hyouk L and Singh K 2020 SwitchTree: in-network computing and traffic analyses with Random Forests *Neural Comput. Appl.* **37** 1–12
- [116] Roy F, Goaz O and Rottenstreich O 2021 Clustreams: data plane clustering *Proc. ACM SIGCOMM Symp. on SDN Research (SOSR)* pp 101–7
- [117] Sanvito D, Siracusano G and Bifulco R 2018 Can the network be the AI accelerator? *Proc. 2018 Workshop on In-Network Computing*
- [118] Zang M, Zheng C, Dittmann L and Zilberman N 2023 Towards continuous threat defense: in-network traffic analysis for IoT gateways *IEEE Int. Things J.* **11** 9244–57
- [119] Zheng C, Rienecker B and Zilberman N 2023 QCOMP: load balancing via in-network reinforcement learning *Proc. 2nd ACM SIGCOMM Workshop on Future of Internet Routing & Addressing* pp 35–40
- [120] Hong X, Zheng C, Zohren S and Zilberman N 2023 LOBIN: in-network machine learning for limit order books *2023 IEEE 24th Int. Conf. on High Performance Switching and Routing (HPSR)* (IEEE) pp 159–66
- [121] Siddique H, Neves M, Kuzniar C and Haque I 2021 Towards network-accelerated ML-based distributed computer vision systems *2021 IEEE 27th Int. Conf. on Parallel and Distributed Systems (ICPADS)* (IEEE) pp 122–9
- [122] Burhanudin U F et al 2021 Light-curve classification with recurrent neural networks for GOTO: dealing with imbalanced data *Mon. Not. R. Astron. Soc.* **505** 4345–61
- [123] Markov I L 2014 Limits on fundamental limits to computation *Nature* **512** 147–54
- [124] Hooker S 2021 The hardware lottery *Commun. ACM* **64** 58–65
- [125] CMS Collaboration 2022 The high-level trigger for the CMS Phase-2 upgrade *PoS ICHEP2022* p 209
- [126] Deng J, Yang Z, Ojima I, Samaras D and Wang F 2022 Artificial intelligence in drug discovery: applications and techniques *Briefings Bioinf.* **23** bbab430
- [127] Bishop B Lawrence livermore national laboratory achieves fusion ignition (available at: www.llnl.gov/archive/news/lawrence-livermore-national-laboratory-achieves-fusion-ignition) (Accessed 11 January 2024)
- [128] Jia Z, Tillman B, Maggioni M and Scarpazza D P 2019 Dissecting the graphcore IPU architecture via microbenchmarking (arXiv:1912.03413)
- [129] Mohan L R, Marshall A, Maddrell-Mander S, O’Hanlon D, Petridis K, Rademacker J, Rege V and Titterton A 2021 Studying the potential of graphcore IPUs for applications in particle physics *Comput. Softw. Big Sci.* **5** 8
- [130] Kita M, Dubiński J, Rokita P and Deja K 2024 Generative diffusion models for fast simulations of particle collisions at CERN (arXiv:2406.03233)
- [131] Qu H, Li C and Qian S 2022 Particle Transformer for Jet Tagging *Int. Conf. on Machine Learning (PMLR)*
- [132] Devlin J, Chang M-W, Lee K and Toutanova K 2019 BERT: pre-training of deep bidirectional transformers for language understanding *North American Chapter of the Association for Computational Linguistics*
- [133] Vig J, Madani A, Varshney L R, Xiong C, Socher R and Fatema Rajani N 2020 BERTology meets biology: interpreting attention in protein language models (arXiv:2006.15222)
- [134] Brandes N, Ofer D, Peleg Y, Rappoport N and Linial M 2022 ProteinBERT: a universal deep-learning model of protein sequence and function *Bioinformatics* **38** 2102–10
- [135] Gopakumar V, Pamela S and Samaddar D 2022 Loss landscape engineering via data regulation on PINNs *Mach. Learn. Appl.* **12** 100464
- [136] Vaswani A et al 2017 Attention is all you need *Advances in Neural Information Processing Systems* vol 30
- [137] Wu T et al 2023 A brief overview of ChatGPT: the history, status quo and potential future development *IEEE/CAA J. Autom. Sin.* **10** 1122–36
- [138] Touvron H et al 2023 Llama: open and efficient foundation language models (arXiv:2302.13971)
- [139] Dodge J et al 2020 Fine-tuning pretrained language models: weight initializations, data orders, and early stopping (arXiv:2002.06305)
- [140] Kaplan J et al 2020 Scaling laws for neural language models (arXiv:2001.08361)
- [141] McKenzie I R et al 2023 Inverse scaling: when bigger isn’t better (arXiv:2306.09479)
- [142] Shoyebi M et al 2019 Megatron-lm: training multi-billion parameter language models using model parallelism (arXiv:1909.08053)
- [143] Shazeer N et al 2017 Outrageously large neural networks: the sparsely-gated mixture-of-experts layer (arXiv:1701.06538)
- [144] Cai T et al 2024 Medusa: simple LLM inference acceleration framework with multiple decoding heads (arXiv:2401.10774)
- [145] Ding J et al 2023 Longnet: scaling transformers to 1,000,000,000 tokens (arXiv:2307.02486)
- [146] Hu E J et al 2021 Lora: low-rank adaptation of large language models (arXiv:2106.09685)

- [147] Lewis P *et al* 2020 Retrieval-augmented generation for knowledge-intensive nlp tasks *Advances in Neural Information Processing Systems* vol 33 pp 9459–74
- [148] Corvoysier David 2023 Make your llama generation time fly with AWS Inferentia2 (available at: <https://huggingface.co/blog/inferentia-llama2>)
- [149] Ray LLMPerf Leaderboard (available at: <https://github.com/ray-project/llmperf-leaderboard>)
- [150] Knight Will 2023 OpenAI's CEO Says the Age of Giant AI Models Is Already Over (available at: www.wired.com/story/openai-ceo-sam-altman-the-age-of-giant-ai-models-is-already-over/)