

Applications of Machine Learning to Lattice Quantum Field Theory

Snowmass 2022 White Paper

Denis Boyda,^{1,2} Salvatore Calì,^{3,2} Sam Foreman,¹ Lena Funcke,^{3,2,4} Daniel C. Hackett,^{3,2,*} Yin Lin,^{3,2} Gert Aarts,^{5,6} Andrei Alexandru,^{7,8} Xiao-Yong Jin,^{1,9} Biagio Lucini,^{10,11} and Phiala E. Shanahan^{3,2,4}

¹*Leadership Computing Facility, Argonne National Laboratory, Argonne, IL 60439, USA*

²*The NSF AI Institute for Artificial Intelligence and Fundamental Interactions*

³*Center for Theoretical Physics, Massachusetts Institute of Technology, Cambridge, MA 02139, USA*

⁴*Co-Design Center for Quantum Advantage (C²QA)*

⁵*Department of Physics, Swansea University, Swansea SA2 8PP, UK*

⁶*European Centre for Theoretical Studies in Nuclear Physics and Related Areas (ECT*)*

⁷*Fondazione Bruno Kessler Strada delle Tabarelle 286, 38123 Villazzano (TN), Italy*

⁸*Physics Department, The George Washington University, Washington, DC 20052, USA*

⁹*Department of Physics, University of Maryland, College Park, MD 20742, USA*

¹⁰*Computational Science Division, Argonne National Laboratory, Argonne, IL 60439, USA*

¹¹*Department of Mathematics, Swansea University, Bay Campus, Swansea SA1 8EN, UK*

¹¹*Swansea Academy of Advanced Computing, Swansea University, Bay Campus, Swansea SA1 8EN, UK*

(Dated: February 15, 2022)

There is great potential to apply machine learning in the area of numerical lattice quantum field theory, but full exploitation of that potential will require new strategies. In this white paper for the Snowmass community planning process, we discuss the unique requirements of machine learning for lattice quantum field theory research and outline what is needed to enable exploration and deployment of this approach in the future.

I. INTRODUCTION

Lattice quantum field theory (LQFT), i.e., QFT discretized on a Euclidean spacetime lattice, is the only framework presently available to perform *ab-initio* QFT calculations with fully controlled, systematically improvable uncertainties when the system of interest exhibits nonperturbative dynamics. Most importantly, this includes quantum chromodynamics (QCD), the component of the Standard Model governing the strong force (see e.g. the recent topical issue of EPJA for a review [1–7]). This important tool of modern physics continues to have great success in providing theory inputs necessary to understand and interpret experimental results. Notably, in quark flavor physics, Lattice QCD (LQCD) provides the hadronic matrix elements necessary to extract the Cabibbo–Kobayashi–Maskawa (CKM) parameters from experimental measurements and hence is critical to tests of CKM unitarity [8]; in precision Higgs physics, LQCD provides the most precise determinations of the strong coupling constant and quark masses that enter in predictions of the branching ratios of the dominant decay modes of the Higgs boson [8]; and lattice studies of QCD thermodynamics are essential to interpret results from relativistic heavy-ion collisions [9].

However, further progress is required to extend these successes of the LQCD approach to meet the needs of ongoing and near-future experimental efforts in high-energy (HEP) and nuclear physics. In particular, while the field is advancing rapidly, for many key applications the necessary LQCD calculations are limited by available computing power; with improvement in algorithms for LQCD, many more important contributions are on the horizon. For example, high-statistics, high-precision LQCD calculations of the hadronic vacuum polarization and hadronic light-by-light scattering contributions to $g - 2$ of the muon could be critical to resolve tension between Standard Model predictions and experiments [10]. At the necessary (sub-percent) level of precision, isospin breaking and QED effects become important, and calculations in the theory with non-degenerate quark masses coupled to electromagnetism are affected by technical and conceptual problems. In addition, so-called disconnected quark diagrams contribute significantly at this level of precision. Stochastic methods have advanced sufficiently to provide access to the otherwise-intractable inverses of Dirac matrices involved in these contributions, but remain expensive and add large additional statistical uncertainties. Further advances in LQFT algorithms will contribute significantly to this pursuit [11].

Similarly, LQCD has the potential to significantly impact the interpretation of the results of long-baseline neutrino experiments by providing constraints on nucleon and (together with effective field theory) nuclear matrix elements involved in neutrino scattering cross sections with heavy nuclei such as ¹²C, water, and ⁴⁰Ar. Precise cross-section

* Editor; dhackett@mit.edu

determination from nuclear models is crucial to reduce the final uncertainties of neutrino parameters, but the precision of these models depends on the uncertainties of their inputs. Many of the relevant matrix elements are difficult to constrain experimentally; in these cases, LQCD is the only viable approach to reduce their uncertainties and meet the needs of experiments [3, 12, 13]. However, direct LQCD calculations of nuclei are computationally expensive and only systems with small atomic numbers are accessible at present. Calculations of larger nuclear systems require new developments to address increasingly severe signal-to-noise problems, which in some cases arise due to numerical sign problems [14], as well as the high combinatoric complexity in Wick contractions. Similarly, computational cost is a significant factor in the computation of nucleon and nuclear matrix elements needed for the interpretation of dark matter direct detection [4] and neutrinoless double beta decay experiments [15, 16], as well as high-precision LQCD results for g_A needed to resolve the outstanding discrepancy between results obtained from different experimental approaches to measuring the neutron lifetime [17]. Extending the study of QCD at finite temperature to nonzero baryon density, as required for nuclear matter, neutron stars, and other phases of dense QCD also leads to a numerical sign problem, due to the complex nature of the Boltzmann weight in the grand-canonical formulation [18, 19].

More broadly, the past decades have seen applications of LQCD across all aspects of hadronic physics. These efforts have already yielded important insights, but in many cases further developments in LQFT technology are necessary to deliver results with the greatest possible impact. For example, the recent development of the quasi- and pseudo-PDF formalisms [20, 21] has enabled LQCD calculations of parton distribution functions (PDFs) and their generalizations, but these methods are still limited by available precision and lattice sizes. Alternately, PDFs may in principle be reconstructed from their moments, which can be computed directly on the lattice; however, issues with power divergent mixings demand presently impractical levels of statistical precision (or novel approaches [22, 23]) for all but the lowest moments. Both of these approaches also involve ill-posed inverse problems, which also appear at finite temperature, in the reconstruction of spectral functions [24], and the extraction of transport coefficients from numerically determined LQCD correlators [25, 26]. Similarly, the correlation functions relevant to computing masses and matrix elements of higher excitations, such as resonances, glueballs, and exotic hadrons suffer from an overwhelming level of numerical noise at the sample sizes accessible with present methods. In some cases, as with nuclear correlation functions, this noise may arise due to a sign problem. Similar concerns apply for scattering amplitudes, which may be reconstructed from the spectra of correlation functions using finite-volume formalisms [27]. Besides QCD, LQFT methods have also been used for direct investigations of strongly coupled models for physics beyond the Standard Model [28–30], such as supersymmetric gauge theories and models where the Higgs boson and/or dark matter are bound states of a new, as-yet unobserved confining force. Models with dilatonic Higgs bosons are particularly computationally demanding. LQFT methods have similarly been applied to address foundational questions in QFT [31–34] and to systems of interest outside of particle physics [35]. To summarize, in many of the cases described above, new physics results can be obtained by more efficient algorithms enabling higher statistical precision, while in others, it is clear that new ideas and novel approaches will be necessary.

From its inception, advancements in LQFT have been achieved by combining increasing computing resources with formal and algorithmic developments. Such developments make new physics targets accessible to the LQFT framework not only by allowing more effective use of existing resources, but also by extending the formalism itself when known approaches are not effective or ill-defined. In this context, emerging machine learning (ML) techniques offer an unprecedented new avenue for exploration, both in terms of increasing efficiency and for new, innovative formulations.

The past few years has seen promising exploratory applications of ML methods to all aspects of LQFT calculations, with many more in progress. To organize these efforts and the discussion below, we divide the workflow of an LQFT calculation into three sequential stages. For concreteness, these stages may be defined by the data types operated upon.

1. *Configuration generation* – Samples of the lattice field degrees of freedom (field configurations) are drawn from the Boltzmann distribution defined by the lattice action. ML applications thus far include novel and/or improved sampling algorithms [36–63] and path-integral contour deformations for finite-density [64–66].
2. *Observable measurement* – Quantities like correlation functions are evaluated over ensembles of field configurations. ML applications thus far include novel methods to extract thermodynamic observables [44], action parameter regression [67], observable approximation [68–70], design of new observables [58, 71–80], and path-integral contour deformations for baryonic correlators [81].
3. *Analysis* – Physically interpretable results are extracted from observable measurements. ML applications thus far include cross-observable regression [82, 83], action parameter regression [67, 84], and new methods for ill-posed inverse problems [85–90].

As discussed further in Sec. II, each of these stages involves different hierarchies of computational scale, requiring different resources and optimizations to apply ML at each stage. Further, each stage has different requirements to maintain full control over uncertainties.

The computational and formal aspects of each of these stages imply specific requirements on applications of ML to

LQFT. In the remainder of this white paper, we first detail key prospects for, and challenges of ML for LQFT. From these, we infer what technical work will be required and discuss what resources the community will need to deploy to enable successful application of ML to LQFT. We close with an outlook on ML for LQFT.

II. PROSPECTS, CHALLENGES, AND REQUIREMENTS OF ML FOR LQFT

A. Prospects

There are important similarities between standard LQFT and ML methods; hence LQFT is already well-positioned to make effective use of emerging ML technologies. We organize the discussion below by these common features, and emphasize how they provide an unusual opportunity for interdisciplinary collaboration and great potential for cross-cutting impact.

Mathematical toolkit. Statistics and linear algebra underpin the standard toolkits of both LQFT and ML. As with particle and condensed matter theory, once a common language is established, researchers from LQFT and those working on other applications of ML find they are concerned with related problems (see e.g. the discussion of symmetry below). By the same analogy, there are likely methods known to one community that can be straightforwardly adapted to solve problems in the other, cf. the development of the HMC algorithm [91] in LQCD and its subsequent adoption in other computational fields, or e.g. using interacting fields as building blocks for neural networks [92]. The potential for mutual benefit offers an unprecedented opportunity for collaboration between the LQCD community and both fields outside physics and industry.

Computing requirements. ML and traditional LQFT methods are computationally intensive. The dominant computational cost in both cases is numerical linear algebra, which is effectively parallelized by GPUs. At present, LQFT requires more tightly interconnected nodes than typical ML applications, but very large models needing fast communications are becoming increasingly common. And, while modern LQFT calculations require expensive double precision (FP64) calculations to avoid unacceptable levels of round-off error, mixed precision algorithms allow efficient LQFT use of hardware with specific optimizations for low-precision formats (e.g. single (FP32), half (FP16), or ML-specific formats (BFLOAT16, TF32)). This approach is already used to accelerate LQCD calculations on GPUs [93] and has been extended to GPU tensor cores [94]; work is ongoing to exploit emerging ML-specialized hardware (e.g. TPUs, IPUs). Thus, hardware that is well-suited for LQCD is likely to be well-suited for ML (and vice versa, increasingly), providing obvious benefits for ease of adoption of ML in LQFT. Given these similarities, ML for LQFT research may lead to faster and better deployment of ML-specialized hardware in traditional LQFT calculations, and computational science generically.

Symmetry. In an ML context, symmetries amount to constraints on a problem which must be learned unless they are incorporated explicitly. Symmetries and invariance/equivariance thus provide a generically useful framework for encoding a priori knowledge about a problem into an ML architecture, which often provides substantially improved performance. Architectures encoding LQFT-relevant symmetries have already been developed in various other contexts. For example, the success of convolutional neural networks in image processing has been substantially driven by their characteristic feature of encoding translational symmetry, which is also a property of LQFTs. Symmetric architectures developed for LQFT will likely be useful for other applications, opening up exciting new possibilities for physics-inspired ML architectures, with a wide range of applications both in academia and industry. For example, rotationally symmetric ML architectures could be useful for applications including 3D modeling, computer vision, and autonomous driving. Symmetric architectures are discussed in more generality in another white paper [95].

Software as a community resource. Significant effort in both the ML and LQFT communities has been dedicated to developing highly optimized open-source codebases. In particular, under collaborative efforts the lattice gauge theory community has created several programming packages (e.g. Grid [96], openQCD [97], QUDA [93], and the SciDAC stack [98]) that enable LQCD calculations to achieve high efficiency on state-of-the-art supercomputers. Similarly, the ML community has developed several independent frameworks, most popularly PyTorch [99], TensorFlow [100], and Jax [101], which allow users to easily utilize GPUs. Another set of frameworks (including horovod [102], DeepSpeed [103], PyTorch DDP [104]) were developed for scaling ML applications to multi-node machines.¹ All of these frameworks are freely available and open-source, and thus may be combined to apply ML to

¹ These techniques rely on efficient communication between multiple workers, but have already demonstrated high performance sustaining more than one exaFLOPS in FP16 precision [105]. However, they may be inefficient for some types of models and can be optimized much further, as in e.g. the DeepSpeed [103] library.

LQFT without incurring the development cost of fully reimplementing either toolkit. Integration of ML and LQFT software may also benefit traditional LQFT calculations. In particular, ML software frameworks are maintained by commercial vendors across a wide range of HPC hardware platforms, which could be leveraged to help address the problem of efficient portability of LQFT codes. Likewise, an increased interest in ML libraries by the LQFT community can further accelerate their development.

B. Challenges

Despite these useful commonalities, there are a number of differences that present challenges for applying ML methods to LQFT problems. By considering them, we can assess what work will be required to bring ML into the standard LQFT toolkit.

Full control over uncertainties. In many ML applications, formal guarantees of exactness are unnecessary and high-quality but approximate solutions are sufficient. Errors can be studied empirically on validation datasets, but characterization of systematic uncertainties may be difficult due to the black-box nature of neural networks.² In contrast, LQFT applications demand correctness, i.e., that all sources of statistical and systematic uncertainty can be estimated reliably. This requirement manifests differently in each stage of the LQFT workflow:

1. *Configuration generation* requires provably exact sampling from the Boltzmann distribution defined by the lattice action, or at least sampling which provides sufficient statistical information to fully correct for violations of exactness. This imposes strict constraints on what ML approaches and architectures are applicable.
2. *Observable measurement* requires control over any violations of asymptotic unbiasedness induced by computing approximations of known observables (e.g. by statistical estimation of a bias correction). Model dependence on training data may present additional concerns for asymptotics. Novel machine-learned observables must be carefully characterized and interpreted with caution.
3. *Analysis* more closely resembles typical ML applications. Most uses of ML at this stage will introduce model dependence, which must be treated as a source of systematic uncertainty and controlled for. This includes dependence on training data, which may induce important model-data correlations. Models with probabilistic interpretations may play a useful role. Estimates of uncertainty are often phrased in a Bayesian context, used in LQFT in e.g. spectral function reconstruction. Many concerns relevant for LQFT applications are discussed in more detail in another white paper on uncertainty quantification [106].

Differentiability. The modern ML toolkit is designed around optimization with stochastic gradient descent, which necessitates that *all* operations be automatically differentiable to enable backpropagation. This is not a typical requirement of LQFT calculations, which more often implement the necessary derivatives of the action manually to maximize performance, and so is not incorporated in any of the major LQFT codebases. Especially for applications to configuration generation and observable measurements, it may be necessary to build automatic differentiation into existing LQFT software, or reimplement parts of the LQFT toolkit inside ML frameworks.

Data hierarchies and computing models. Typical ML applications operate on large volumes of data points, each of size \sim KBs. Although this hierarchy applies in the analysis stage of LQFT calculations, it is very dissimilar to the relevant scales in state-of-the-art LQCD configuration generation and observable measurements, which involve processing relatively small volumes of field configurations and propagators of size \sim GBs – 100s of GBs. These large data sizes mean applying ML methods at state-of-the-art scale will require a higher degree of model parallelism than typical for ML, where data parallelism alone is often sufficient, and where the need for model parallelism is more often due to large model sizes.³ Although the recent successes of very large ML models has led to software support for model parallelism, this infrastructure is still in its infancy. Engineering will be required to adapt ML software to make efficient use of high-performance computing resources. Further, problem-specific optimizations are more important for model parallel schemes, which clashes with the problem-agnostic approach employed by ML software frameworks. Applying ML to LQFT may require specialized software incorporating LQFT-specific optimizations.

² Some recent activities in ML focus on out-of-distribution detection.

³ In a data parallel scheme, parallel workers each apply the entire ML model to multiple pieces of data. This scheme involves relatively little communication between workers, requiring only the synchronization of gradients of model parameters once per backwards pass. In a model parallel scheme, the work of applying an ML model to a single piece of data is divided between multiple workers. This may be accomplished e.g. by dividing the problem across sublattices, although other schemes like pipeline parallelism are possible.

C. Requirements

From these concerns, we conclude that successful application of ML techniques in lattice QCD will require significant efforts in two primary directions:

1. *Exploratory research* – The particular requirements of LQFT mean that out-of-the-box ML solutions often are not directly applicable. Applying ML to LQFT while retaining full control over uncertainties will require research and development of novel constructions with appropriate properties. Given the breadth of ML methods already available, and with the rapid ongoing development of ML, this amounts to a need for substantial exploratory research. The work required is experimental and inherently computationally demanding, especially when testing scalability on large systems.
2. *Software engineering* – Following precedent, using ML-based methods at state-of-the-art scale will require development and long-term maintenance of publicly available software infrastructure by the LQFT community. This software must be portable and highly optimized for efficient use of available hardware resources, as well as well-tested, well-documented, and open source to ensure scientific validity and verifiability. Development can be accelerated by leveraging the large amount of work and domain expertise incorporated into existing LQFT and ML software frameworks. However, at present, these pieces of infrastructure are entirely separate, and as discussed above have been designed with different considerations (cf. differentiability, data hierarchies and computing models). Delivering high performance in ML for LQFT at scale thus represents a substantial software engineering task, which will require experts with extensive knowledge in both domains.

III. STRATEGIES TO ENABLE ML FOR LQFT

Computing strategies. New policies for allocating computing time are needed to support exploratory algorithm development and applications of ML-based approaches. Usually, allocations are requested for a fixed amount of computing time, to run a specific set of computations on a predetermined schedule, and are granted based on what physics results the proposed computations are projected to enable. This does not match with the nature of exploratory ML research, which involves an iterative experimental process with each experiment guiding what computations are run next. Iteration timescales are typically of order days or weeks, much less than year-long computing allocations. More generally, even given a production-ready approach, present allocation policies are incompatible with any computation that involves training a model: until the training is actually carried out, the precise cost and outcome will not be known to high accuracy. Employing all talent available, and especially avoiding unfairly shutting out researchers from smaller and less-funded institutions, will require making computing resources broadly available.

Community standards and resources. As discussed above, ML for LQFT will require development and maintenance of specialized software toolkits. Just as with software, trained models should be treated as a community resource, particularly for at-scale applications where training may be expensive. Best practices must be established for what information is required to constitute a verifiable result [107], likely involving sharing of code and models alongside publications. To these ends, centralized infrastructure will be needed to enable sharing of trained models and the code to use them.

Career paths. Permanent positions in ML for physics must be made available to support the exploratory research and software engineering described above. This includes both traditional academic jobs and positions for specialists in technical roles, such as Research Software Engineers (RSEs). Increased support for research scientists working on ML for physics will promote building a vibrant interdisciplinary community and bridge gaps between different subfields of physics and beyond. Given the utility of ML outside academia, these concerns are especially important to retain talented early-career researchers, particularly those engaged in valuable but highly technical work. Policies and practices should be adjusted for a definition of physics research and education that incorporates computer science and applied mathematics relevant to ML. Cross-disciplinary collaborations may provide rapid progress as well as access to substantial non-traditional funding and computational resources, including from industrial partners. Hiring and graduate admissions should consider applicants from outside physics; it may be that the best-prepared candidate for ML-based research is a student with an undergraduate degree in computer science, or a researcher from another field or even industry. This applies especially in education, as discussed in greater detail in another white paper [108]. While specialized degrees can play a useful role, ML for physics is a legitimate topic for a physics degree. Qualifying exam practices should allow for students specialized in computation. Physics departments should incorporate computing and ML in their curricula; besides supporting the development of ML for the physics community, this will benefit the majority of students who go on to careers outside research.

IV. OUTLOOK

From vector machines to BlueGene to GPUs, the LQFT community has a long history of leadership in rapidly adopting new computing technologies, and driving their development as they emerge. The community now has the opportunity to position itself for similar leadership in ML. Promising proof-of-principle results across every aspect of the LQFT workflow and rapidly growing engagement in work at this intersection both illustrate potential to deliver transformative advances in the immediate future.

Realizing this potential will require intentional investment of human and computing resources by the community in ways distinct from those that have driven traditional algorithms research. Bringing ML to bear will require extensive, computationally demanding exploratory research incorporating developments from the broader field of ML, requiring updates to the way computing resources are allocated. When promising methods are identified, deploying them at state-of-the-art scale will require software engineering to meet the existing high standards for performance.

The development and deployment of novel ML algorithms for LQFT has great potential for cross-cutting impact. Besides its widespread adoption in industry, ML methods have already been applied to various problems in particle/nuclear physics [109–111], astronomy [112, 113], condensed-matter physics [114], computational fluid dynamics [115, 116], quantum chemistry [117], and many further fields of physics and beyond. As outlined in this white paper, applications to LQFT bring specific challenges to the forefront, such as symmetries and the need for provably exact algorithms. The solutions to these challenges may provide opportunities not only for science but industrial applications as well. Moreover, most of these considerations would be equally beneficial for other emerging computational research directions, including tensor networks and quantum computing.

Developing the interdisciplinary workforce who will carry out this work will require investment and advocacy by the LQFT community. For example, early career researchers are responsible for much of the ongoing research in ML for LQFT; the community risks losing this talent pool if it cannot provide permanent positions to retain them, either by opening physics positions to researchers with such an algorithmic focus, by opening computer science positions to researchers with a focus on physics problems, or by creating permanent positions at the interdisciplinary boundaries. This talent pool can be further expanded through the inclusion of computing and ML in physics curricula, and more generally opening the field of physics to admit the specialization necessary for ML-based research. These same concerns apply not only for LQFT, but for ML for physics in general. Nevertheless, if these challenges can be successfully negotiated, the interdisciplinary nature and generalizability of ML methods provide a unique opportunity for collaboration between fields and with industry, opening the doors to new intellectual, computational, and funding resources which may have significant impact on the state-of-the-art in HEP theory.

ACKNOWLEDGMENTS

DB, LF, DCH, YL, SC, and PES are supported in part by the U.S. Department of Energy, Office of Science, Office of Nuclear Physics, under grant Contract Number DE-SC0011090. PES is additionally supported by the National Science Foundation under EAGER grant 2035015, by the U.S. DOE Early Career Award DE-SC0021006, by a NEC research award, and by the Carl G and Shirley Sontheimer Research Fund. YL is also supported by the National Science Foundation award PHY-2019786. LF is also supported by the U.S. Department of Energy, Office of Science, National Quantum Information Science Research Centers, Co-design Center for Quantum Advantage (C²QA) under contract number DE-SC0012704, by the DOE QuantiSED Consortium under subcontract number 675352, by the National Science Foundation under Cooperative Agreement PHY-2019786 (The NSF AI Institute for Artificial Intelligence and Fundamental Interactions, <http://iaifi.org/>), and by the U.S. Department of Energy, Office of Science, Office of Nuclear Physics under grant contract number DE-SC0021006. DB, SF and XJ were supported by the Argonne Leadership Computing Facility, which is a U.S. Department of Energy Office of Science User Facility operated under contract DE-AC02-06CH11357. SF and XJ were additionally supported by the Exascale Computing Project (17-SC-20-SC), a collaborative effort of the U.S. Department of Energy Office of Science and the National Nuclear Security Administration. GA and BL are supported in part by the UKRI Science and Technology Facilities Council (STFC) Consolidated Grant ST/T000813/1. The work of BL is further supported in part by the Royal Society Wolfson Research Merit Award WM170010 and by the Leverhulme Foundation Research Fellowship RF-2020-461/9. AA is supported in part by U.S. DOE Grant No. DE-FG02-95ER40907.

[1] Richard C. Brower, Anna Hasenfratz, Ethan T. Neil, Simon Catterall, George Fleming, Joel Giedt, Enrico Rinaldi, David Schaich, Evan Weinberg, and Oliver Witzel (USQCD), “Lattice Gauge Theory for Physics Beyond the Standard Model,”

- Eur. Phys. J. A* **55**, 198 (2019), [arXiv:1904.09964 \[hep-lat\]](#).
- [2] Christoph Lehner *et al.* (USQCD), “Opportunities for Lattice QCD in Quark and Lepton Flavor Physics,” *Eur. Phys. J. A* **55**, 195 (2019), [arXiv:1904.09479 \[hep-lat\]](#).
 - [3] Andreas S. Kronfeld, David G. Richards, William Detmold, Rajan Gupta, Huey-Wen Lin, Keh-Fei Liu, Aaron S. Meyer, Raza Sufian, and Sergey Syritsyn (USQCD), “Lattice QCD and Neutrino-Nucleus Scattering,” *Eur. Phys. J. A* **55**, 196 (2019), [arXiv:1904.09931 \[hep-lat\]](#).
 - [4] Vincenzo Cirigliano, Zohreh Davoudi, Tanmoy Bhattacharya, Taku Izubuchi, Phiala E. Shanahan, Sergey Syritsyn, and Michael L. Wagman (USQCD), “The Role of Lattice QCD in Searches for Violations of Fundamental Symmetries and Signals for New Physics,” *Eur. Phys. J. A* **55**, 197 (2019), [arXiv:1904.09704 \[hep-lat\]](#).
 - [5] William Detmold, Robert G. Edwards, Jozef J. Dudek, Michael Engelhardt, Huey-Wen Lin, Stefan Meinel, Kostas Orginos, and Phiala Shanahan (USQCD), “Hadrons and Nuclei,” *Eur. Phys. J. A* **55**, 193 (2019), [arXiv:1904.09512 \[hep-lat\]](#).
 - [6] Alexei Bazavov, Frithjof Karsch, Swagato Mukherjee, and Peter Petreczky (USQCD), “Hot-dense Lattice QCD: USQCD whitepaper 2018,” *Eur. Phys. J. A* **55**, 194 (2019), [arXiv:1904.09951 \[hep-lat\]](#).
 - [7] Bálint Joó, Chulwoo Jung, Norman H. Christ, William Detmold, Robert Edwards, Martin Savage, and Phiala Shanahan (USQCD), “Status and Future Perspectives for Lattice Gauge Theory Calculations to the Exascale and Beyond,” *Eur. Phys. J. A* **55**, 199 (2019), [arXiv:1904.09725 \[hep-lat\]](#).
 - [8] Y. Aoki *et al.*, “FLAG Review 2021,” (2021), [arXiv:2111.09849 \[hep-lat\]](#).
 - [9] Claudia Ratti, “Lattice QCD and heavy ion collisions: a review of recent progress,” *Rept. Prog. Phys.* **81**, 084301 (2018), [arXiv:1804.07810 \[hep-lat\]](#).
 - [10] B. Abi *et al.* (Muon $g-2$), “Measurement of the Positive Muon Anomalous Magnetic Moment to 0.46 ppm,” *Phys. Rev. Lett.* **126**, 141801 (2021), [arXiv:2104.03281 \[hep-ex\]](#).
 - [11] A. X. El-Khadra, “Review of muon $g-2$,” (2021), The 38th International Symposium on Lattice Field Theory.
 - [12] Aaron S. Meyer, André Walker-Loud, and Callum Wilkinson, “Status of Lattice QCD Determination of Nucleon Form Factors and their Relevance for the Few-GeV Neutrino Program,” (2022), [arXiv:2201.01839 \[hep-lat\]](#).
 - [13] L. Alvarez-Ruso *et al.* (NuSTEC), “NuSTEC White Paper: Status and challenges of neutrino-nucleus scattering,” *Prog. Part. Nucl. Phys.* **100**, 1–68 (2018), [arXiv:1706.03621 \[hep-ph\]](#).
 - [14] Michael L. Wagman and Martin J. Savage, “Statistics of baryon correlation functions in lattice QCD,” *Phys. Rev. D* **96**, 114508 (2017), [arXiv:1611.07643 \[hep-lat\]](#).
 - [15] Jonathan Engel and Javier Menéndez, “Status and Future of Nuclear Matrix Elements for Neutrinoless Double-Beta Decay: A Review,” *Rept. Prog. Phys.* **80**, 046301 (2017), [arXiv:1610.06548 \[nucl-th\]](#).
 - [16] Michelle J. Dolinski, Alan W. P. Poon, and Werner Rodejohann, “Neutrinoless Double-Beta Decay: Status and Prospects,” *Ann. Rev. Nucl. Part. Sci.* **69**, 219–251 (2019), [arXiv:1902.04097 \[nucl-ex\]](#).
 - [17] Martín González-Alonso, Oscar Naviliat-Cuncic, and Nathal Severijns, “New physics searches in nuclear and neutron β decay,” *Prog. Part. Nucl. Phys.* **104**, 165–223 (2019), [arXiv:1803.08732 \[hep-ph\]](#).
 - [18] Philippe de Forcrand, “Simulating QCD at finite density,” *PoS LAT2009*, 010 (2009), [arXiv:1005.0539 \[hep-lat\]](#).
 - [19] Gert Aarts, “Introductory lectures on lattice QCD at nonzero baryon number,” *J. Phys. Conf. Ser.* **706**, 022004 (2016), [arXiv:1512.05145 \[hep-lat\]](#).
 - [20] Xiangdong Ji, “Parton Physics on a Euclidean Lattice,” *Phys. Rev. Lett.* **110**, 262002 (2013), [arXiv:1305.1539 \[hep-ph\]](#).
 - [21] A. V. Radyushkin, “Quasi-parton distribution functions, momentum distributions, and pseudo-parton distribution functions,” *Phys. Rev. D* **96**, 034025 (2017), [arXiv:1705.01488 \[hep-ph\]](#).
 - [22] A. J. Chambers, R. Horsley, Y. Nakamura, H. Perlt, P. E. L. Rakow, G. Schierholz, A. Schiller, K. Somfleth, R. D. Young, and J. M. Zanotti, “Nucleon Structure Functions from Operator Product Expansion on the Lattice,” *Phys. Rev. Lett.* **118**, 242001 (2017), [arXiv:1703.01153 \[hep-lat\]](#).
 - [23] Zohreh Davoudi and Martin J. Savage, “Restoration of Rotational Symmetry in the Continuum Limit of Lattice Field Theories,” *Phys. Rev. D* **86**, 054505 (2012), [arXiv:1204.4146 \[hep-lat\]](#).
 - [24] M. Asakawa, T. Hatsuda, and Y. Nakahara, “Maximum entropy analysis of the spectral functions in lattice QCD,” *Prog. Part. Nucl. Phys.* **46**, 459–508 (2001), [arXiv:hep-lat/0011040](#).
 - [25] Gert Aarts and Jose Maria Martínez Resco, “Transport coefficients, spectral functions and the lattice,” *JHEP* **04**, 053 (2002), [arXiv:hep-ph/0203177](#).
 - [26] Harvey B. Meyer, “Transport Properties of the Quark-Gluon Plasma: A Lattice QCD Perspective,” *Eur. Phys. J. A* **47**, 86 (2011), [arXiv:1104.3708 \[hep-lat\]](#).
 - [27] Martin Luscher, “Trivializing maps, the Wilson flow and the HMC algorithm,” *Commun. Math. Phys.* **293**, 899–919 (2010), [arXiv:0907.5491 \[hep-lat\]](#).
 - [28] Thomas DeGrand, “Lattice tests of beyond standard model dynamics,” *Rev. Mod. Phys.* **88**, 015001 (2016), [arXiv:1510.05018 \[hep-ph\]](#).
 - [29] V. Drach, “Composite electroweak sectors on the lattice,” *PoS LATTICE2019*, 242 (2020), [arXiv:2005.01002 \[hep-lat\]](#).
 - [30] David Schaich, “Progress and prospects of lattice supersymmetry,” *PoS LATTICE2018*, 005 (2019), [arXiv:1810.09282 \[hep-lat\]](#).
 - [31] Pilar Hernández and Fernando Romero-López, “The large N_c limit of QCD on the lattice,” *Eur. Phys. J. A* **57**, 52 (2021), [arXiv:2012.03331 \[hep-lat\]](#).
 - [32] Biagio Lucini and Marco Panero, “SU(N) gauge theories at large N,” *Phys. Rept.* **526**, 93–163 (2013), [arXiv:1210.4997 \[hep-th\]](#).
 - [33] J. Greensite, “The Confinement problem in lattice gauge theory,” *Prog. Part. Nucl. Phys.* **51**, 1 (2003), [arXiv:hep-lat/0301023](#).

- [34] Georges Ripka, *Dual superconductor models of color confinement*, Vol. 639 (2004) [arXiv:hep-ph/0310102](#).
- [35] Manu Mathur and T. P. Sreeraj, “Lattice Gauge Theories and Spin Models,” *Phys. Rev. D* **94**, 085029 (2016), [arXiv:1604.00315 \[hep-lat\]](#).
- [36] Lei Wang, “Exploring cluster Monte Carlo updates with Boltzmann machines,” *Phys. Rev. E* **96**, 051301 (2017).
- [37] Li Huang and Lei Wang, “Accelerated Monte Carlo simulations with restricted Boltzmann machines,” *Physical Review B* **95**, – (2017).
- [38] Jiaming Song, Shengjia Zhao, and Stefano Ermon, “A-NICE-MC: Adversarial training for MCMC,” in *Advances in Neural Information Processing Systems* (2017) pp. 5140–5150.
- [39] Shuo-Hui Li and Lei Wang, “Neural Network Renormalization Group,” *Phys. Rev. Lett.* **121**, 260601 (2018).
- [40] Daniel Levy, Matthew D. Hoffman, and Jascha Sohl-Dickstein, “Generalizing hamiltonian monte carlo with neural networks,” (2018), [arXiv:1711.09268 \[stat.ML\]](#).
- [41] M. S. Albergo, G. Kanwar, and P. E. Shanahan, “Flow-based generative models for Markov chain Monte Carlo in lattice field theory,” *Phys. Rev. D* **100**, 034515 (2019), [arXiv:1904.12072 \[hep-lat\]](#).
- [42] Dian Wu, Lei Wang, and Pan Zhang, “Solving Statistical Mechanics Using Variational Autoregressive Networks,” *Phys. Rev. Lett.* **122**, 080602 (2019).
- [43] Kim A. Nicoli, Shinichi Nakajima, Nils Strodthoff, Wojciech Samek, Klaus-Robert Müller, and Pan Kessel, “Asymptotically unbiased estimation of physical observables with neural samplers,” *Phys. Rev. E* **101**, 023304 (2020), [arXiv:1910.13496 \[cond-mat.stat-mech\]](#).
- [44] Kim A. Nicoli, Christopher J. Anders, Lena Funcke, Tobias Hartung, Karl Jansen, Pan Kessel, Shinichi Nakajima, and Paolo Stornati, “Estimation of Thermodynamic Observables in Lattice Field Theories with Deep Generative Models,” *Phys. Rev. Lett.* **126**, 032001 (2021), [arXiv:2007.07115 \[hep-lat\]](#).
- [45] Jan M. Pawłowski and Julian M. Urban, “Reducing Autocorrelation Times in Lattice Simulations with Generative Adversarial Networks,” *Mach. Learn. Sci. Tech.* **1**, 045011 (2020), [arXiv:1811.03533 \[hep-lat\]](#).
- [46] Gurtej Kanwar, Michael S. Albergo, Denis Boyda, Kyle Cranmer, Daniel C. Hackett, Sébastien Racanière, Danilo Jimenez Rezende, and Phiala E. Shanahan, “Equivariant flow-based sampling for lattice gauge theory,” *Phys. Rev. Lett.* **125**, 121601 (2020), [arXiv:2003.06413 \[hep-lat\]](#).
- [47] Denis Boyda, Gurtej Kanwar, Sébastien Racanière, Danilo Jimenez Rezende, Michael S. Albergo, Kyle Cranmer, Daniel C. Hackett, and Phiala E. Shanahan, “Sampling using $SU(N)$ gauge equivariant flows,” *Phys. Rev. D* **103**, 074504 (2021), [arXiv:2008.05456 \[hep-lat\]](#).
- [48] Michael S. Albergo, Gurtej Kanwar, Sébastien Racanière, Danilo J. Rezende, Julian M. Urban, Denis Boyda, Kyle Cranmer, Daniel C. Hackett, and Phiala E. Shanahan, “Flow-based sampling for fermionic lattice field theories,” (2021), [arXiv:2106.05934 \[hep-lat\]](#).
- [49] Daniel C. Hackett, Chung-Chun Hsieh, Michael S. Albergo, Denis Boyda, Jiunn-Wei Chen, Kai-Feng Chen, Kyle Cranmer, Gurtej Kanwar, and Phiala E. Shanahan, “Flow-based sampling for multimodal distributions in lattice field theory,” (2021), [arXiv:2107.00734 \[hep-lat\]](#).
- [50] Marylou Gabrié, Grant M. Rotskoff, and Eric Vanden-Eijnden, “Adaptive Monte Carlo augmented with normalizing flows,” (2021), [arXiv:2105.12603 \[physics.data-an\]](#).
- [51] Luigi Del Debbio, Joe Marsh Rossney, and Michael Wilson, “Efficient Modelling of Trivializing Maps for Lattice ϕ^4 Theory Using Normalizing Flows: A First Look at Scalability,” (2021), [arXiv:2105.12481 \[hep-lat\]](#).
- [52] Dian Wu, Riccardo Rossi, and Giuseppe Carleo, “Unbiased Monte Carlo Cluster Updates with Autoregressive Neural Networks,” (2021), [arXiv:2105.05650 \[cond-mat.stat-mech\]](#).
- [53] Sam Foreman, Xiao-Yong Jin, and James C. Osborn, “Deep Learning Hamiltonian Monte Carlo,” in *9th International Conference on Learning Representations* (2021) [arXiv:2105.03418 \[hep-lat\]](#).
- [54] Sam Foreman, Taku Izubuchi, Luchang Jin, Xiao-Yong Jin, James C. Osborn, and Akio Tomiya, “HMC with Normalizing Flows,” in *38th International Symposium on Lattice Field Theory* (2021) [arXiv:2112.01586 \[cs.LG\]](#).
- [55] Sam Foreman, Xiao-Yong Jin, and James C. Osborn, “LeapfrogLayers: A Trainable Framework for Effective Topological Sampling,” in *38th International Symposium on Lattice Field Theory* (2021) [arXiv:2112.01582 \[hep-lat\]](#).
- [56] Xiao-Yong Jin, “Neural Network Field Transformation and Its Application in HMC,” in *38th International Symposium on Lattice Field Theory* (2022) [arXiv:2201.01862 \[hep-lat\]](#).
- [57] Jacob Finkenrath, “Tackling critical slowing down using global correction steps with equivariant flows: the case of the Schwinger model,” (2022), [arXiv:2201.02216 \[hep-lat\]](#).
- [58] Dimitrios Bachtis, Gert Aarts, and Biagio Lucini, “Extending machine learning classification capabilities with histogram reweighting,” *Phys. Rev. E* **102**, 033303 (2020), [arXiv:2004.14341 \[cond-mat.stat-mech\]](#).
- [59] Pim de Haan, Corrado Rainone, Miranda C. N. Cheng, and Roberto Bondesan, “Scaling Up Machine Learning For Quantum Field Theory with Equivariant Continuous Flows,” (2021), [arXiv:2110.02673 \[cs.LG\]](#).
- [60] Guido Cossu, Luigi Del Debbio, Tommaso Giani, Ava Khamseh, and Michael Wilson, “Machine learning determination of dynamical parameters: The Ising model case,” *Phys. Rev. B* **100**, 064304 (2019), [arXiv:1810.11503 \[physics.comp-ph\]](#).
- [61] Akinori Tanaka and Akio Tomiya, “Towards reduction of autocorrelation in HMC by machine learning,” (2017), [arXiv:1712.03893 \[hep-lat\]](#).
- [62] Akio Tomiya and Yuki Nagai, “Gauge covariant neural network for 4 dimensional non-abelian gauge theory,” (2021), [arXiv:2103.11965 \[hep-lat\]](#).
- [63] Dimitrios Bachtis, Gert Aarts, Francesco Di Renzo, and Biagio Lucini, “Inverse renormalization group in quantum field theory,” *Phys. Rev. Lett.* (to appear) (2021), [arXiv:2107.00466 \[hep-lat\]](#).
- [64] Andrei Alexandru, Paulo F. Bedaque, Henry Lamm, and Scott Lawrence, “Deep Learning Beyond Lefschetz Thimbles,”

- Phys. Rev. **D96**, 094505 (2017).
- [65] Scott Lawrence and Yukari Yamauchi, “Normalizing Flows and the Real-Time Sign Problem,” *Phys. Rev. D* **103**, 114509 (2021), [arXiv:2101.05755 \[hep-lat\]](#).
- [66] Jan-Lukas Wynen, Evan Berkowitz, Stefan Krieg, Thomas Luu, and Johann Ostmeyer, “Machine learning to alleviate Hubbard-model sign problems,” *Phys. Rev. B* **103**, 125153 (2021), [arXiv:2006.11221 \[cond-mat.str-el\]](#).
- [67] Phiala E. Shanahan, Daniel Trewartha, and William Detmold, “Machine learning action parameters in lattice quantum chromodynamics,” *Phys. Rev. D* **97**, 094506 (2018), [arXiv:1801.05784 \[hep-lat\]](#).
- [68] Matteo Favoni, Andreas Ipp, David I. Müller, and Daniel Schuh, “Lattice Gauge Equivariant Convolutional Neural Networks,” *Phys. Rev. Lett.* **128**, 032003 (2022), [arXiv:2012.12901 \[hep-lat\]](#).
- [69] Srinath Bulusu, Matteo Favoni, Andreas Ipp, David I. Müller, and Daniel Schuh, “Generalization capabilities of translationally equivariant neural networks,” *Phys. Rev. D* **104**, 074504 (2021), [arXiv:2103.14686 \[hep-lat\]](#).
- [70] Takuya Matsumoto, Masakiyo Kitazawa, and Yasuhiro Kohno, “Classifying topological charge in SU(3) Yang–Mills theory with machine learning,” *PTEP* **2021**, 023D01 (2021), [arXiv:1909.06238 \[hep-lat\]](#).
- [71] D. L. Boyda, M. N. Chernodub, N. V. Gerasimeniuk, V. A. Goy, S. D. Liubimov, and A. V. Molochkov, “Finding the deconfinement temperature in lattice Yang–Mills theories from outside the scaling window with machine learning,” *Phys. Rev. D* **103**, 014509 (2021), [arXiv:2009.10971 \[hep-lat\]](#).
- [72] Dimitrios Bachtis, Gert Aarts, and Biagio Lucini, “Adding machine learning within Hamiltonians: Renormalization group transformations, symmetry breaking and restoration,” *Phys. Rev. Res.* **3**, 013134 (2021), [arXiv:2010.00054 \[hep-lat\]](#).
- [73] Andrea Palermo, Lucio Anderlini, Maria Paola Lombardo, Andrey Kotov, and Anton Trunin, “Machine learning approaches to the QCD transition,” in *38th International Symposium on Lattice Field Theory* (2021) [arXiv:2111.05216 \[hep-lat\]](#).
- [74] D. R. Tan, J. H. Peng, Y. H. Tseng, and F. J. Jiang, “A universal neural network for learning phases,” *Eur. Phys. J. Plus* **136**, 1116 (2021), [arXiv:2103.10846 \[cond-mat.dis-nn\]](#).
- [75] Chian-De Li, Deng-Ruei Tan, and Fu-Jiun Jiang, “Applications of neural networks to the studies of phase transitions of two-dimensional Potts models,” *Annals Phys.* **391**, 312–331 (2018), [arXiv:1703.02369 \[cond-mat.dis-nn\]](#).
- [76] Sebastian Johann Wetzels and Manuel Scherzer, “Machine Learning of Explicit Order Parameters: From the Ising Model to SU(2) Lattice Gauge Theory,” *Phys. Rev. B* **96**, 184410 (2017), [arXiv:1705.05582 \[cond-mat.stat-mech\]](#).
- [77] Constantia Alexandrou, Andreas Athenodorou, Charalambos Chrysostomou, and Srijit Paul, “The critical temperature of the 2D-Ising model through Deep Learning Autoencoders,” *Eur. Phys. J. B* **93**, 226 (2020), [arXiv:1903.03506 \[cond-mat.stat-mech\]](#).
- [78] Stefan Blücher, Lukas Kades, Jan M. Pawłowski, Nils Strodthoff, and Julian M. Urban, “Towards novel insights in lattice field theory with explainable machine learning,” *Phys. Rev. D* **101**, 094507 (2020), [arXiv:2003.01504 \[hep-lat\]](#).
- [79] Hon Man Yau and Nan Su, “On the generalizability of artificial neural networks in spin models,” (2020), [arXiv:2006.15021 \[cond-mat.dis-nn\]](#).
- [80] Kai Zhou, Gergely Endrődi, Long-Gang Pang, and Horst Stöcker, “Regressive and generative neural networks for scalar field theory,” *Phys. Rev. D* **100**, 011501 (2019), [arXiv:1810.12879 \[hep-lat\]](#).
- [81] William Detmold, Gurtej Kanwar, Henry Lamm, Michael L. Wagman, and Neill C. Warrington, “Path integral contour deformations for observables in $SU(N)$ gauge theory,” *Phys. Rev. D* **103**, 094517 (2021), [arXiv:2101.12668 \[hep-lat\]](#).
- [82] Boram Yoon, Tanmoy Bhattacharya, and Rajan Gupta, “Machine Learning Estimators for Lattice QCD Observables,” *Phys. Rev. D* **100**, 014504 (2019), [arXiv:1807.05971 \[hep-lat\]](#).
- [83] Rui Zhang, Zhouyou Fan, Ruizi Li, Huey-Wen Lin, and Boram Yoon, “Machine-learning prediction for quasiparton distribution function matrix elements,” *Phys. Rev. D* **101**, 034516 (2020), [arXiv:1909.10990 \[hep-lat\]](#).
- [84] R. J. Hudspith and D. Mohler, “A fully non-perturbative charm-quark tuning using machine learning,” (2021), [arXiv:2112.01997 \[hep-lat\]](#).
- [85] Lukas Kades, Jan M. Pawłowski, Alexander Rothkopf, Manuel Scherzer, Julian M. Urban, Sebastian J. Wetzels, Nicolas Wink, and Felix P. G. Ziegler, “Spectral Reconstruction with Deep Neural Networks,” *Phys. Rev. D* **102**, 096001 (2020), [arXiv:1905.04305 \[physics.comp-ph\]](#).
- [86] Sam Offler, Gert Aarts, Chris Allton, Jonas Glesaaen, Benjamin Jäger, Seyong Kim, Maria Paola Lombardo, Sinead M. Ryan, and Jon-Ivar Skullerud, “News from bottomonium spectral functions in thermal QCD,” *PoS LATTICE2019*, 076 (2019), [arXiv:1912.12900 \[hep-lat\]](#).
- [87] Jan Horak, Jan M. Pawłowski, José Rodríguez-Quintero, Jonas Turnwald, Julian M. Urban, Nicolas Wink, and Savvas Zafeiropoulos, “Reconstructing QCD Spectral Functions with Gaussian Processes,” (2021), [arXiv:2107.13464 \[hep-ph\]](#).
- [88] S. Y. Chen, H. T. Ding, F. Y. Liu, G. Papp, and C. B. Yang, “Machine learning spectral functions in lattice QCD,” (2021), [arXiv:2110.13521 \[hep-lat\]](#).
- [89] Lingxiao Wang, Shuzhe Shi, and Kai Zhou, “Automatic differentiation approach for reconstructing spectral functions with neural networks,” in *35th Conference on Neural Information Processing Systems* (2021) [arXiv:2112.06206 \[physics.comp-ph\]](#).
- [90] Shuzhe Shi, Lingxiao Wang, and Kai Zhou, “Rethinking the ill-posedness of the spectral function reconstruction - why is it fundamentally hard and how Artificial Neural Networks can help,” (2022), [arXiv:2201.02564 \[hep-ph\]](#).
- [91] S. Duane, A. D. Kennedy, B. J. Pendleton, and D. Roweth, “Hybrid Monte Carlo,” *Phys. Lett. B* **195**, 216–222 (1987).
- [92] Dimitrios Bachtis, Gert Aarts, and Biagio Lucini, “Quantum field-theoretic machine learning,” *Phys. Rev. D* **103**, 074510 (2021), [arXiv:2102.09449 \[hep-lat\]](#).
- [93] M. A. Clark, R. Babich, K. Barros, R. C. Brower, and C. Rebbi, “Solving Lattice QCD systems of equations using mixed precision solvers on GPUs,” *Comput. Phys. Commun.* **181**, 1517–1528 (2010), [arXiv:0911.3191 \[hep-lat\]](#).

- [94] Jiqun Tu, M. A. Clark, Chulwoo Jung, and Robert Mawhinney, “Solving DWF Dirac Equation Using Multi-splitting Preconditioned Conjugate Gradient with Tensor Cores on NVIDIA GPUs,” (2021), [10.1145/3468267.3470613](https://arxiv.org/abs/2104.05615), [arXiv:2104.05615](https://arxiv.org/abs/2104.05615) [hep-lat].
- [95] “Symmetry Group Equivariant Architectures for Physics,” in *Snowmass 2022*, in preparation.
- [96] Peter Boyle, Azusa Yamaguchi, Guido Cossu, and Antonin Portelli, “Grid: A next generation data parallel C++ QCD library,” (2015), [arXiv:1512.03487](https://arxiv.org/abs/1512.03487) [hep-lat].
- [97] M. Lüscher *et al.*, “openQCD,” <https://luscher.web.cern.ch/luscher/openQCD/> (2022).
- [98] R. Edwards *et al.* (USQCD), “SciDAC,” <https://usqcd-software.github.io/> (2022).
- [99] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala, “Pytorch: An imperative style, high-performance deep learning library,” in *Advances in Neural Information Processing Systems 32*, edited by H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (Curran Associates, Inc., 2019) pp. 8024–8035.
- [100] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dandelion Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng, “TensorFlow: Large-scale machine learning on heterogeneous systems,” (2015), software available from tensorflow.org.
- [101] James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang, “JAX: composable transformations of Python+NumPy programs,” (2018).
- [102] Alexander Sergeev and Mike Del Balso, “Horovod: fast and easy distributed deep learning in TensorFlow,” arXiv preprint [arXiv:1802.05799](https://arxiv.org/abs/1802.05799) (2018).
- [103] Jeff Rasley, Samyam Rajbhandari, Olatunji Ruwase, and Yuxiong He, “Deepspeed: System optimizations enable training deep learning models with over 100 billion parameters,” in *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining* (Association for Computing Machinery, New York, NY, USA, 2020) p. 3505–3506.
- [104] Shen Li, Yanli Zhao, Rohan Varma, Omkar Salpekar, Pieter Noordhuis, Teng Li, Adam Paszke, Jeff Smith, Brian Vaughan, Pritam Damania, and Soumith Chintala, “Pytorch distributed: Experiences on accelerating data parallel training,” (2020), [arXiv:2006.15704](https://arxiv.org/abs/2006.15704) [cs.DC].
- [105] Nouamane Laanait, Joshua Romero, Junqi Yin, M Todd Young, Sean Treichler, Vitalii Starchenko, Albina Borisevich, Alex Sergeev, and Michael Matheson, “Exascale deep learning for scientific inverse problems,” arXiv preprint [arXiv:1909.11150](https://arxiv.org/abs/1909.11150) (2019).
- [106] “Uncertainty Quantification in Machine Learning,” in *Snowmass 2022*, in preparation.
- [107] Mark D Wilkinson, Michel Dumontier, IJsbrand Jan Aalbersberg, Gabrielle Appleton, Myles Axton, Arie Baak, Niklas Blomberg, Jan-Willem Boiten, Luiz Bonino da Silva Santos, Philip E Bourne, *et al.*, “The fair guiding principles for scientific data management and stewardship,” *Scientific data* **3**, 1–9 (2016).
- [108] “Data Science & Machine Learning Education in HEP,” in *Snowmass 2022*, in preparation.
- [109] Matthew Feickert and Benjamin Nachman, “A Living Review of Machine Learning for Particle Physics,” (2021), [arXiv:2102.02770](https://arxiv.org/abs/2102.02770) [hep-ph].
- [110] Amber Boehnlein *et al.*, “Artificial Intelligence and Machine Learning in Nuclear Physics,” (2021), [arXiv:2112.02309](https://arxiv.org/abs/2112.02309) [nucl-th].
- [111] Akinori Tanaka, Akio Tomiya, and Kōji Hashimoto, *Deep Learning and Physics* (Springer, 2021).
- [112] Elena Cuoco *et al.*, “Enhancing Gravitational-Wave Science with Machine Learning,” *Mach. Learn. Sci. Tech.* **2**, 011002 (2021), [arXiv:2005.03745](https://arxiv.org/abs/2005.03745) [astro-ph.HE].
- [113] Dalya Baron, “Machine learning in astronomy: a practical overview,” (2019), [arXiv:1904.07248](https://arxiv.org/abs/1904.07248) [astro-ph.IM].
- [114] Edwin A. Bedolla-Montiel, Luis Carlos Padierna, and Ramón Castañeda Priego, “Machine Learning for Condensed Matter Physics,” *J. Phys. Condens. Matter* **33**, 053001 (2021), [arXiv:2005.14228](https://arxiv.org/abs/2005.14228) [physics.comp-ph].
- [115] Steven L. Brunton, Bernd R. Noack, and Petros Koumoutsakos, “Machine learning for fluid mechanics,” *Annual Review of Fluid Mechanics* **52**, 477–508 (2020), <https://doi.org/10.1146/annurev-fluid-010719-060214>.
- [116] Dmitrii Kochkov, Jamie A. Smith, Ayya Alieva, Qing Wang, Michael P. Brenner, and Stephan Hoyer, “Machine learning-accelerated computational fluid dynamics,” *Proceedings of the National Academy of Sciences* **118** (2021), [10.1073/pnas.2101784118](https://doi.org/10.1073/pnas.2101784118), <https://www.pnas.org/content/118/21/e2101784118.full.pdf>.
- [117] Kristof T Schütt, Stefan Chmiela, O Anatole von Lilienfeld, Alexandre Tkatchenko, Koji Tsuda, and Klaus-Robert Müller, “Machine learning meets quantum physics,” *Lecture Notes in Physics* (2020).