nature reviews physics
https://doi.org/10.1038/s42254-024-00798-x

Perspective

Check for updates

Physics-driven learning for inverse problems in quantum chromodynamics

Gert Aarts [1], Kenji Fukushima [2], Tetsuo Hatsuda [3], Andreas Ipp [4], Shuzhe Shi [5], Lingxiao Wang [3] ✉ & Kai Zhou [6,7]

Abstract

The integration of deep learning techniques and physics-driven designs is reforming the way we address inverse problems, in which accurate physical properties are extracted from complex observations. This is particularly relevant for quantum chromodynamics (QCD) – the theory of strong interactions – with its inherent challenges in interpreting observational data and demanding computational approaches. This Perspective highlights advances of physics-driven learning methods, focusing on predictions of physical quantities towards QCD physics and drawing connections to machine learning. Physics-driven learning can extract quantities from data more efficiently in a probabilistic framework because embedding priors can reduce the optimization effort. In the application of first-principles lattice QCD calculations and QCD physics of hadrons, neutron stars and heavy-ion collisions, we focus on learning physically relevant quantities, such as perfect actions, spectral functions, hadron interactions, equations of state and nuclear structure. We also emphasize the potential of physics-driven designs of generative models beyond QCD physics.

Sections

[1]Department of Physics, Swansea University, Swansea, UK. [2]Department of Physics, The University of Tokyo, Tokyo, Japan. [3]Interdisciplinary Theoretical and Mathematical Sciences Program (iTHEMS), RIKEN, Wako, Japan. [4]Institute for Theoretical Physics, TU Wien, Vienna, Austria. [5]Department of Physics, Tsinghua University, Beijing, China. [6]School of Science and Engineering, The Chinese University of Hong Kong, Shenzhen, Guangdong, China. [7]Frankfurt Institute for Advanced Studies, Frankfurt am Main, Germany. ✉e-mail: lingxiao.wang@riken.jp

Key points

- Inverse problems in physical sciences determine causes or parameters from observations.
- Physics-driven learning integrates domain-specific physical knowledge into machine learning to solve inverse problems.
- Physics-driven learning can help to extract physical properties and build probability distributions from data.
- In quantum chromodynamics, physics-driven learning can deduce hadron forces, dense matter equations of state, and nuclear structure.
- Physics-driven designs can innovate the development of deep learning and generative models.

Introduction

Quantum chromodynamics (QCD) is the fundamental theory describing the strong interactions of quarks and gluons, which make up hadrons such as protons and neutrons[1]. Characterized by its non-perturbative nature at low energies, QCD physics derives phenomena such as confinement (permanent binding of quarks and gluons within hadrons) and asymptotic freedom (weaker quark interactions at high energies). QCD physics is crucial for understanding nuclear and extreme matter; first-principle lattice QCD calculations[2], compact star observations[3] and relativistic nuclear collision experiments[4] are approaches promising to further this understanding.

Exploring QCD physics and decoding its associated phenomena involve many challenging inverse problems[5,6], which are needed to determine causes or parameters from consequent observations. Unlike forward problems, which predict outcomes from known factors, inverse problems start from results to do reverse engineering[7]. They are essential in fields wherein direct measurements are impractical, such as medical imaging, geophysics and astrophysics. Inverse problems in QCD physics involve identifying strong interaction properties from intricate measurements. Examples include extracting hadron spectral functions[8] from lattice observables, reconstructing dense matter equations of state (EoSs) from compact star observations[9], and identifying quark–gluon plasma (QGP) properties from heavy-ion collision (HIC) experiments[6].

Machine learning (ML) techniques, as a modern branch of artificial intelligence (AI), are becoming increasingly important in QCD physics[6,10], providing tools to identify intricate patterns and extract structures from complex data sets[11,12]. Bayesian inference can deduce causal parameters from uncertain observations, and deep models are trained to learn physical properties from well-prepared data. Recently, advanced developments in physics-driven learning have integrated physical priors into machine learning explicitly[6,13,14], thereby improving deep models to produce physically interpretable and accurate results. This Perspective aims to build a bridge between solving inverse problems in physics and advancing machine learning techniques. For QCD physicists, we present the latest advances and applications of physics-driven learning in areas such as first-principle lattice calculations, hadron physics, neutron stars and heavy-ion collisions. We emphasize that the more interpretable machine learning techniques will benefit the inverse problem solutions. Meanwhile, the applications of physics-driven learning are not limited to QCD physics. For machine learning experts and other physicists interested in inverse problems, this Perspective also presents a concise plan of how to embed physics knowledge into the learning procedure.

Physics-driven learning

Physical concepts have been fundamental in the development of ML[15]. Mathematical frameworks and problem-solving techniques in physics have inspired algorithms and models of deep learning, including energy-based models[16], maximum entropy thermodynamics[17], non-equilibrium stochastic dynamics[18] and physics-inspired neural network design and optimization[19]. The influential physics-driven learning paradigm is helpful both in solving inverse problems and in feeding back AI innovations, such as developing more physically interpretable generative models and aligning deep models, to the physical world.

In solving inverse problems for QCD physics, Bayesian inference and deep learning have emerged as indispensable tools, each contributing uniquely to the resolution of the involved challenges.

Bayesian inference provides a rigorous probabilistic framework for understanding and quantifying uncertainties inherent in QCD physics. Using Bayesian statistics, one systematically incorporates prior knowledge to update beliefs on the considered physics based on observations and the posterior distribution $p(\theta|D)$ of the parameters θ. The observations D have to be obtained using Bayes' theorem[20], $p(\theta|D) = p(D|\theta)p(\theta)/p(D)$, where $p(D|\theta)$ is the likelihood, $p(\theta)$ is the prior and $p(D)$ is the evidence. This allows the integration of prior knowledge about the physics, providing a coherent method to update the causal parameters of the studied physics as new data become available, which also enables the estimation of parameter uncertainties and the construction of credible intervals.

As a branch of ML, deep learning[11,12] offers a complementary solution for inverse problems, with its capacity to model complex and non-linear relationships using deep neural networks. In the context of QCD, deep learning can be used to approximate the inverse mapping from observations to the underlying physics. This is achieved by training deep models on data sets generated from theoretical model simulations or experimental measurements. The network learns to infer the latent parameters θ from observations D, performing inverse engineering, which usually works under the principle of maximum likelihood estimation. With this objective, one can choose different architectures of artificial neural networks (ANNs), such as convolutional neural networks (CNNs)[21], recurrent neural networks[22], residual networks[23], graph neural networks[24] and transformers[25] for various inverse problems.

However, Bayesian inference is often stalled by the need to optimize a large number of parameters, making it computationally expensive and impractical for identifying detailed information within physical quantities. Deep learning models, conversely, are typically designed for data-driven tasks and lack the explicit physical constraints, making them less suitable for accurately decoding physical quantities. This mismatch highlights the need for new methods tailored to the accurate requirements of physical systems. Meanwhile, solving inverse problems involves inferring causes from observations, which can be particularly difficult when data are incomplete or noisy. These problems are often ill-posed or ill-conditioned, which means that solutions may not exist, may not be unique or may not be stable[7]. Regularization techniques are routinely used to obtain meaningful solutions.

Integrating physics priors into deep learning methods is a promising strategy to address these challenges simultaneously and effectively[6,26]. Designing deep models explicitly with physics knowledge can introduce physics regularization, reduce the parameter space and implement efficient gradient-based optimization methods. The deep models can serve as universal approximators[12] to represent the underlying physics quantities, learning from observations within a Bayesian perspective and optimizing within auto-differential processes described by physical equations. This emerging paradigm is called physics-driven deep learning, and its goal is to construct better solutions of inverse problems. The physical knowledge that can be integrated includes symmetries, physics principles, well-developed physics equations, and physical data from simulations or experiments, as shown in Fig. 1.

- Symmetry, a cornerstone of modern physics, can be used to improve learning performance explicitly. In general, embedding symmetries introduces a scheme for sharing parameters in deep models[11], thus reducing the number of parameters and preventing over-fitting[27]. Symmetries, such as translational, rotational and permutation invariance, can be inherently incorporated into network architectures including CNNs[12,28], Euclidean neural networks[29] and graph neural networks[30]. In QCD physics, it is particularly important to embed gauge invariance and equivariance in deep models[31–33].
- Physics principles, such as causality, continuity, positive definiteness and asymptotic behaviour, are also crucial in ensuring that solutions are physically meaningful. These properties can be implemented into model designs – either through customized loss functions or specialized activation functions[34,35] – to ensure that the output respects these fundamental constraints.
- Physics equations, in particular differential equations governing systems, can serve as priors. These laws provide essential constraints that guide the learning process, ensuring that solutions adhere to physical realities. In contrast to the physics-informed neural networks[36], in which physical equations are always introduced as additional regularization terms, physical equations can be explicitly coded into the optimization of deep learning. For instance, ordinary or partial differential equations are automatically differentiable[37] and can, therefore, be encoded into the forward process of a deep-learning model, and their gradients can be computed in reverse optimizations for training. Overall, this approach leads to more robust and physically controllable outcomes[14].
- Physical data, whether obtained from experiments or simulations[38], serve as a form of regularization, which helps to align the model outputs with physical truths[39]. In particular, for ill-posed inverse problems and those that require initial verification of the existence of the inverse mapping, such a regularization not only can ground the learning process in reality but can also mitigate the risk of obtaining non-physical solutions.

Developing deep learning models with specific physics knowledge can further enhance their capability and effectiveness[15,40]. For example, domain-specific insights can be incorporated into the design of deep generative models whose backbone is the inverse modelling of underlying probability distributions[12]. Whether starting from exact likelihood estimation or not, one can approximate the underlying distributions in data and generate reliable samples. These models can benefit from existing and well-verified physics rules.

QCD physics

In this section, we provide examples of physics-driven learning in the context of QCD in four areas of interest: lattice QCD, hadron physics, neutron stars and heavy-ion collisions.

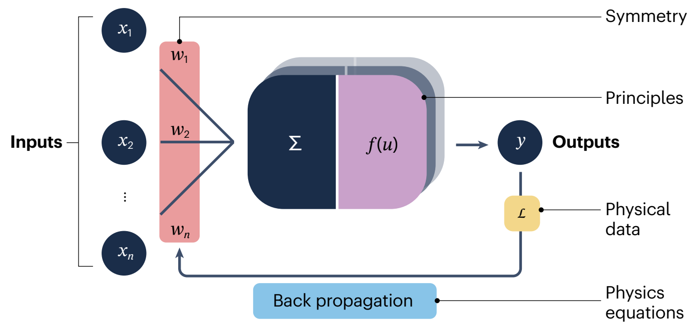

Fig. 1 | Physics-driven designs for deep learning. In a deep neural network model, the weights $\{w\} = w_1, w_2, \ldots, w_n$ connect the inputs $\{x\} = x_1, x_2, \ldots, x_n$ and the outputs y with summation Σ and non-linear activation functions $f(u)$. In a single layer, the equation can be written as $y = f(\Sigma_{i=1}^{n} x_i w_i)$. Symmetries can be encoded by sharing weights $\{w\}$, and other principles can be represented by different activation functions $f(u)$. Owing to its differentiable properties, physics equations can be explicitly utilized in the back propagation algorithm within an automatic differentiable framework[37]. Physical data provide guidance for the outputs from deep models when computing the loss functions.

Lattice QCD

Lattice QCD (LQCD) provides a first-principle, non-perturbative approach to study strong interactions[2]. Progress over the past 50 years has been driven by a combined improvement of algorithms and an increase in computational power, executing simulations on the largest supercomputers. The workhorse of LQCD is the Markov Chain Monte Carlo algorithm, to generate large ensembles of field configurations in four-dimensional space-time. However, this algorithm cannot access the full QCD phase diagram; the low-temperature, moderate-to-high-density region is still out of bounds owing to the sign problem: at nonzero baryon chemical potential, the quark determinant is complex and the Monte Carlo-based importance sampling method cannot be used[41].

ML provides a new tool to investigate and enhance LQCD simulations, and indeed ML is already used to study many aspects of LQCD, including configuration generation and observable measurement and analysis[42]. Many applications are still exploratory and most are developed in theories that are easier to solve than QCD; nevertheless, there is promise and we will discuss selected examples here. An important aspect is the precision and accuracy of the solutions: to compete with or improve on well-established methods, it is necessary to demonstrate that ML-driven algorithms[33] can deliver high-precision results with controllable systematic and statistical uncertainties.

Before starting an LQCD simulation, one needs to determine parameters in the lattice action. To minimize discretization effects, one can use perfect or fixed-point actions, which are classically free from discretization errors but, in principle, come with an infinite number of couplings[43,44]. In pioneering works, ML has been used to learn action parameters[45] with the aim of improving lattice simulations[46] and gaining explainable insight into lattice systems[47]. Generically, such networks had to learn gauge symmetry, or it was included through manually selected small Wilson loop structures. A more scalable approach to learning fixed-point actions has recently been emerged[48,49] by using lattice gauge-equivariant CNNs (L-CNNs) to determine parametrizations, which have been shown to be superior to the ones considered previously. L-CNNs are essential in achieving this improved performance,

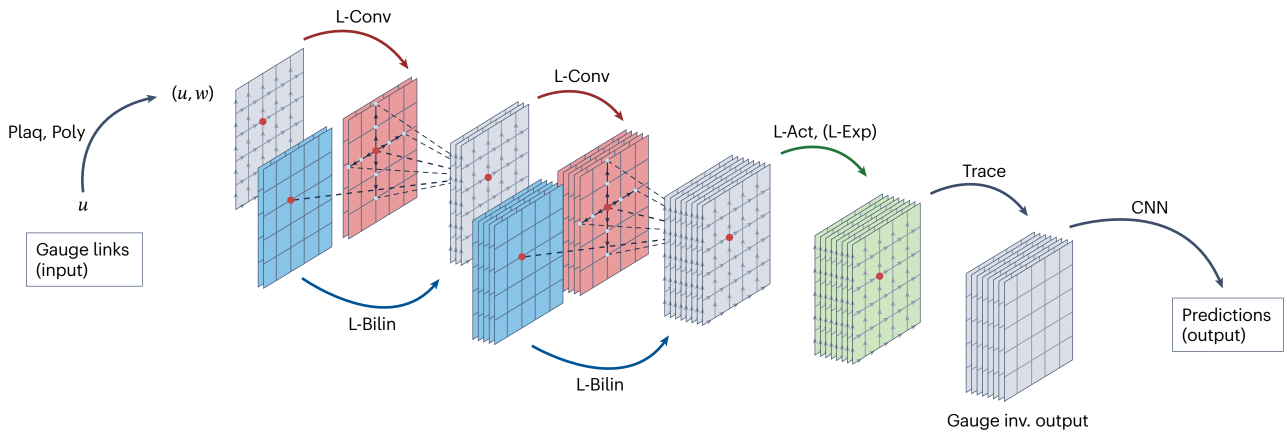

Fig. 2 | A lattice gauge-equivariant convolutional neural network (L-CNN). This architecture processes data defined on a lattice, representing quantum field theory in discretized space-time. Each layer in the L-CNN is carefully constructed to preserve gauge symmetry, a property crucial for ensuring physically consistent predictions in quantum field theory. The network first processes the input lattice data given by gauge links u using simple Wilson loops on the lattice such as plaquettes (Plaq) or Polyakov loops (Poly) to create objects w that transform locally. These objects are then combined into progressively more complex Wilson loops while maintaining the symmetry through specialized convolutional (L-Conv) and bilinear (L-Bilin) operations. Additional gauge equivariant activation functions (L-Act) or exponentiation layers (L-Exp) can modify the local fields in a gauge equivariant manner. Finally, the network generates gauge-invariant (inv.) outputs through a trace layer (Trace) that can be processed by standard convolutional layers to produce the desired physical predictions. Unlike conventional convolutional neural networks (CNNs), this design is robust to random and adversarial gauge transformations, making it essential for simulations of fundamental physics. Reprinted from ref. 32, CC BY 4.0.

as Fig. 2 shows, because they provide[32] a very general formulation of a lattice action, for example, as functions of arbitrarily sized Wilson loops, while maintaining exact gauge symmetry in the network architecture. Similar to the universal approximation theorem for neural networks, L-CNNs can in principle approximate any gauge-covariant function on the lattice, thus offering a powerful tool for capturing a wide range of gauge-invariant features. This generality allows L-CNNs to be adapted to different physical problems by tailoring the network architecture to specific tasks, such as action parametrization, energy minimization or even generative modelling of gauge field configurations.

Owing to its link with generative AI, ensemble generation has received most of the attention. Early applications[50,51] of this approach used generative adversarial networks to generate field configurations in a two-dimensional scalar field theory. By now, normalizing flow is the best developed approach and has been reviewed extensively elsewhere[33,52]. Another way to perform ensemble generation is based on diffusion models, and they are of interest for a number of reasons: the underlying idea comes from physics[18,53] and lattice field theorists find stochastic updates easy to understand. Indeed, the relation between diffusion models and stochastic quantization[54] – a well-known technique to generate quantum field configurations from a stochastic process in a fictitious time dimension – has been pointed out recently[55], and first applications to scalar and U(1) gauge field theories in two dimensions have been implemented[55–57]. Unlike the normalizing flow approach, diffusion models learn from configurations that were previously generated with any alternative approach, such as hybrid Monte Carlo methods, but the trained diffusion model can subsequently be incorporated in the Markov Chain to increase the size of ensembles. Further connections of diffusion models to the path integral[58] and renormalization group flows[59] have been pointed out as well. Despite the computational efficiency of ensemble generation, exactness is particularly important if this method is to compete with well-established Monte Carlo-based methods. Additional accept–reject procedures[33] and importance sampling estimates[60] can both help to improve exactness.

Because of the prominence of the renormalization group in QCD, exploring it from an inverse perspective is a natural choice. In this machine learning context, the idea is for the inverse renormalization group, originally proposed for spin systems[61], to learn a transformation that undoes a standard renormalization group transformation by using transposed convolutions[62]. If the inverse transformation is local, it can be applied over and over again, generating ever larger lattices that get closer to criticality. This idea was first applied to scalar fields in two dimensions[62] and also to a hard-to-simulate disordered system – the three-dimensional Edwards–Anderson model[63], which describes spin glasses with randomly interacting spins on a three-dimensional lattice. The crucial step to make this work is to design a 'good' renormalization group step, which requires understanding of the physics in the critical regime.

An exciting area wherein machine learning and lattice field theory (LFT) intersect lies in their shared ability to describe systems with many fluctuating degrees of freedom. The two approaches share notable similarities that can be mutually beneficial, even though their starting points differ widely: QCD is a fundamental theory in which the probability distribution (or action) is dictated by symmetries and renormalization conditions, whereas ML models typically learn or approximate distributions from data. One example in which concepts in ML and LFT intersect nontrivially is the development of gauge-equivariant networks, in which local (gauge) symmetries, which are essential in lattice systems[32,64–66], are respected by ML architectures[67,68]. Moreover, ML is increasingly being used to learn order parameters and determine the phase structure in simulated theories[69], which has been successfully extended to lattice gauge theories[70,71]. There is also growing interest in analysing deep models through the lens of (quantum) field theory, particularly building on the relationship between deep neural networks in the infinite width limit and Gaussian processes or free fields[72,73]. Specifically, LFT concepts can be integrated into deep models by adding LFT interactions to nodes[74], or LFT analyses can provide new perspectives, as demonstrated in the case of Gaussian restricted Boltzmann machines[75] or via the relation to Dyson Brownian motion and random matrix theory[76]. We expect the synthesis of LFT insights with ML frameworks to advance both fields and to deepen the understanding of complex systems.

Hadron physics

Much understanding of the subatomic world is rooted in hadron physics. Recently, machine learning techniques have become useful tools to further this understanding as they help to understand hadron spectra, search for exotic hadrons and study hadron interactions.

The theoretical extraction of hadron spectra from LQCD, which is an inverse problem, is a challenging task because there is only a finite amount of LQCD data with statistical noise. Apart from the standard method of extracting hadron masses using long-range temporal correlations, the maximum entropy method[8,77], based on Bayesian inference, provides a robust framework for the extraction of hadron spectral functions from LQCD data using information entropy for regularization. Recently, automatic differentiation methods (Fig. 3) based on maximum likelihood estimation have been developed. The spectral function is first represented by a neural network ansatz, then the difference between the predictions and the real observations is computed as a loss function $\mathscr{L}$, whose minimization can be back-propagated to optimize parameters $\{\theta\}$ in neural networks, as $\partial\mathscr{L}/\partial\theta$. In the forward process, the physical integral is explicitly encoded as a sum to compute the predictions. The principles that the spectral function has to be continuous and positive are also explicitly incorporated into the flexible neural network representations[35,78].

The exploration of exotic hadrons has become an exciting frontier in hadron physics[79]. Exotic states, including tetraquarks, pentaquarks and other multiquark states, challenge the traditional meson and baryon picture. Candidates of exotic hadrons are often found experimentally as peaks near certain decay thresholds, and it is important to identify theoretically whether they correspond to bound, virtual or resonant poles in the complex plane. Bound poles represent stable systems in which particles are permanently bound, virtual poles reflect transient interactions that affect scattering without forming stable states, and resonant poles represent unstable, short-lived states that decay into other particles. One can detect such poles from partial wave amplitudes using deep neural networks[80,81] by taking into account the general properties of the S-matrix, such as analyticity and unitarity. ML studies on the properties of the narrow pentaquark $P_c(4312)$[82] and other exotic states[83] are currently an active field or research, and neural network wavefunctions for hadrons[84,85] may help to identify and characterize these exotic states.

Hadronic interactions have a crucial role not only in the investigation of the structure of exotic hadrons but also in the understanding of how atomic nuclei form and of the physics inside neutron stars[3]. One approach to study hadron interactions is the HAL QCD method[86–88], which has been proposed as a way to build effective potentials between hadrons from their spatial correlations (equal-time Nambu–Bethe–Salpeter amplitude) measured on the lattice. This approach could bridge the gap between LQCD theory and experimental data[89,90]. Deriving the potential from spatial correlations in LQCD is also an inverse problem with which physics-driven learning can help: symmetric neural networks, trained on LQCD data with constraints of asymptotic behaviours, can model general non-local potentials in the Schrödinger equation[91,92].

Neutron stars

One of the unsolved questions in modern physics is the state of baryonic matter in highly compressed environments, such as those achieved in heavy-ion collisions and occurring in neutron stars[93] with their high densities. The typical observed values of neutron star masses range

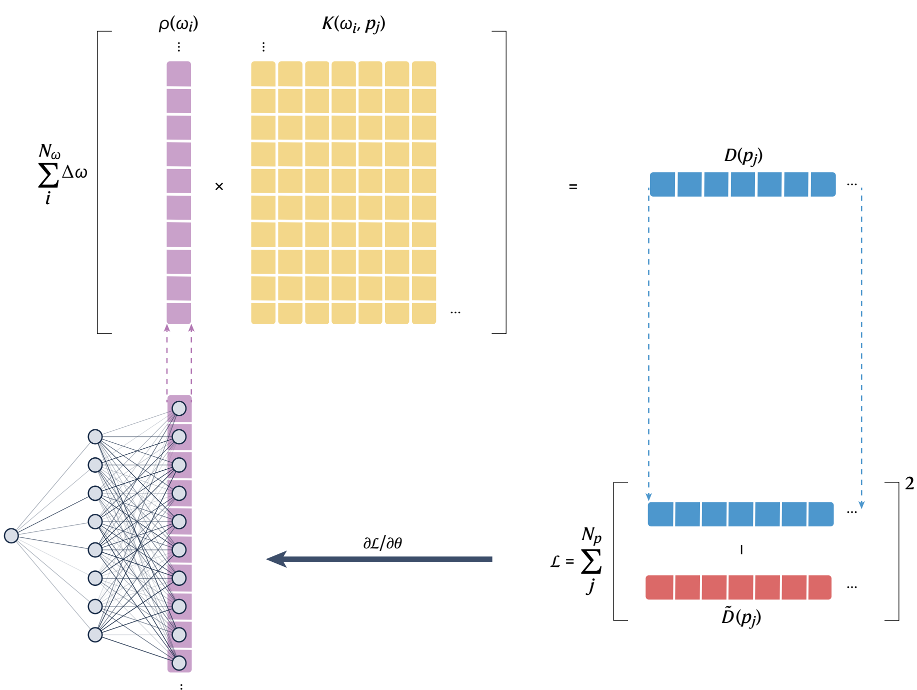

Fig. 3 | Automatic differential framework to reconstruct hadron spectral functions from observations[78]. Neural networks have outputs as a list representation of spectrum $\rho(\omega_i)$. The convolutional operation between $\rho(\omega_i)$ and kernel function $K(\omega,p)$ gives the predicted observations $D(p)$ as $D(p_j)=\sum_i^{N_\omega}\Delta\omega\rho(\omega_i)K(\omega_i,p_j)$. The difference between real observations and predicted ones is used to compute the loss function $\mathscr{L}=\sum_j^{N_p}(D(p_j)-\tilde{D}(p_j))^2$ for optimizing the weights $\{\theta\}$ of neural networks with gradient $\partial\mathscr{L}/\partial\theta$. The activation functions of the neural network can be set, for example, as Softplus, to meet the positive definition principle. Adapted from ref. 78, CC BY 4.0.

Glossary

Artificial neural networks
(ANNs). Models inspired by the structure and function of biological neural networks in human brains.

Deep neural networks
Complex ANNs with multiple layers, including input, output and at least one hidden layer.

Convolutional neural networks
(CNNs). Excel with image, speech and audio inputs; they consist of three main types of layers: convolutional, pooling and fully connected layers.

Recurrent neural networks
Bi-directional ANNs, unlike the uni-directional feedforward network; they allow outputs from nodes to influence subsequent inputs to the same nodes.

from 1.4 to 2 times the solar mass with radii 11–13 km, corresponding to a maximum density of up to 10^{12} kg cm^{-3}. Identifying the internal structures of such neutron stars from theoretical predictions and experimental observations is crucial to unravel the nature of high-density QCD[3,9]. Because General Relativity (GR) and QCD are well-established theories, the bulk properties of the neutron stars should be uniquely solvable as a well-posed problem (see Fig. 4). Nevertheless, knowing the theories that lead to the correct answer is different from knowing the answer itself. In the case of neutron stars, a relation between the pressure and the energy density, called the equation of state, is essential to the hydrostatic equations that describe the balance between the inward gravitational force and the outward pressure derivative. Although Monte Carlo simulations of QCD on the lattice can effectively explore EoSs at finite temperatures and zero baryon densities, the sign problem introduces large uncertainties as baryon densities increase, preventing their direct application to neutron star EoS.

For a given QCD EoS, one can deal with equations in GR to uniquely fix the distribution of the masses and the radii of the neutron stars, ensuring that the mass–radius relation in principle draws a single curve. Conversely, one can reconstruct the QCD EoS from the experimentally observed mass–radius relation if it is well constrained, as shown in Fig. 4. However, both the quality and quantity of observed data is insufficient, making this typical inverse problem ill-posed. To nevertheless facilitate a solution, one can estimate the most probable candidate among all EoS possibilities. A naive strategy along these lines would be the following: generate random EoSs with an assumed prior distribution, weigh the likelihood to get observations from these EoSs, and subsequently quantify the posterior distribution[94]. These procedures are systematized in a form of Bayesian analysis[95,96].

Alternatively, deep neural networks can be trained in a supervised manner with inference[97], and there are proposals to use a hybrid of Bayesian inference and neural networks[98,99] to reduce EoSs. In addition, to embed physics priors more explicitly, neural network-represented EoSs have been implemented in the automatic differentiation framework[100,101], into which the hydrostatic equations are coded. Other physics-driven ideas to address this inverse problem include physical data augmentation[102] and the redesign of activation functions at the output layer of neural network-represented EoSs to satisfy the principle of causality and microscopic stability condition[34].

Heavy-ion collisions

Relativistic heavy-ion collisions provide the unique chance to create and explore the extreme state of QCD matter in terrestrial laboratories, such as the Relativistic Heavy-Ion Collider and the Large Hadron Collider. In HICs, charged ions are accelerated to nearly the speed of light, collide and produce thousands of hadrons, photons and leptons, which are tracked by the detectors. Analyses of the final-state particles indicate that a new state of matter – the quark–gluon plasma – may be formed in the early stage of the collision. The formation of the QGP is an important test of the quark model and QCD theory because it happens in the colour deconfinement phase, during which quarks and gluons can travel distances that greatly exceed the size of hadrons. Soon after their formation in HICs, QGP droplets drastically expand – driven by the pressure gradient – cool down, and form hadrons. An inevitable challenge in HICs is inversely revealing the underlying QCD physics that determines the highly complex and fast-evolving collision dynamics. Owing to the interdependent influence on different experimental observables, state-of-the-art theoretical simulations for the collision dynamics involve many physical uncertainties[103], including fluctuations of the initial state and bulk and transport properties of the QCD matter. Deep learning techniques, assisted by physics prior, have proven useful in addressing such inverse problems in HICs.

A central task of HICs is to determine the QCD EoS and the relevant phase structure. Recently, deep learning models, specifically CNNs, have been used to directly map the final-state particle spectra in HICs back to the bulk properties of QCD[104]. As shown in Fig. 5, this study explored how important features of the QCD EoS, such as the order of the QCD phase transitions, can persist through the complex evolution of the system during the collision. Despite the intricate nature of relativistic hydrodynamic processes, CNNs were able to reliably

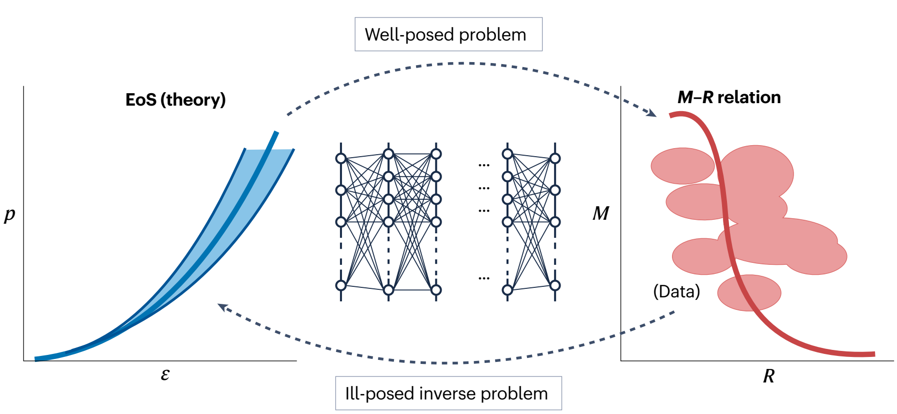

Fig. 4 | The ill-posed inverse problem from observational data to theory in the context of neutron star physics. At the theoretical level, mapping between the equation of state (EoS) and the mass–radius (M–R) relation is well-posed, but the inferred probable EoS has a probability distribution reflecting the quality and quantity of data. Left panel shows the EoS, which defines the relationship between pressure (P) and energy density (ϵ) inside neutron stars. The right panel represents the corresponding mass–radius data. In the middle, a neural network can be used to learn the inverse mapping from the M–R data back to the EoS[97] or represent the forward process for inverse inference[100].

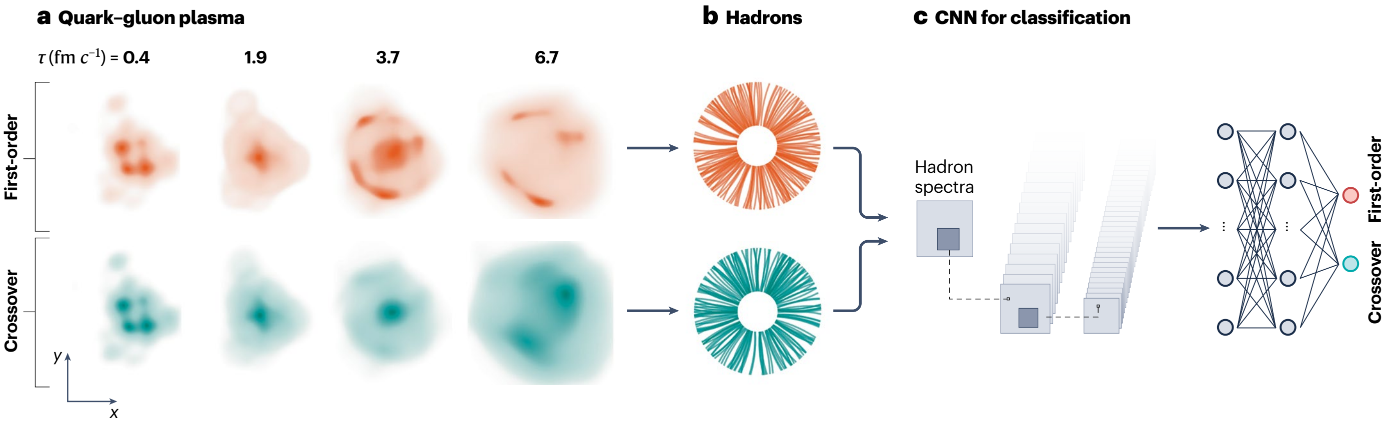

Fig. 5 | Quantum chromodynamics (QCD) transition binary classification using convolutional neural networks (CNNs) with final particle spectra from heavy-ion collisions as input. a, The system starts with the creation of a quark–gluon plasma (QGP) in the collision. The hydrodynamic evolution over time (τ) of the QGP encodes different types of QCD phase transitions, such as crossover and first-order transitions. **b**, After the QGP cools down, it transitions into a hadronic phase. The final-state particles, including various types of hadrons, are detected experimentally, providing the observable data. **c**, CNNs are used to perform binary classification by identifying signals of the QCD phase transition from the hadron spectra. The spectra are preprocessed into image-like data, which makes CNNs a suitable tool to extract features and to classify the observations as either a crossover or a first-order phase transition.

decode this information from the final-state spectra of pions, even when accounting for uncertainties such as initial state fluctuations and shear viscosity effects. This way of analysing inverse HICs was further developed to tackle non-equilibrium phase transition scenarios and afterburner hadronic cascades[105–107]. It also allowed the inclusion of experimental detector effects by taking the hits and tracks of particles from the detector in the form of point cloud – defined as an unordered list of points with their attributes such as momentum and charge – as direct input[108].

The extraction of the QCD EoS has also benefited from physics-driven Bayesian analyses, in which knowledge is encoded in the priors and the model emulator. Inference of the temperature-dependent speed of sound from high-energy HICs[109] provided the first experimental evidence of a smooth transition between the QGP and the hadronic phase at zero baryon density. Bayesian analysis was successfully applied to extract the density-dependent QCD EoS[110] and bulk and shear viscosities[111–113]. It is interesting to note that there is good consistency between such data-driven inference of viscosities and those computed from a quasi-particle model with temperature-dependent masses, extracted from a physics-driven deep learning method to match the EoSs calculated with LQCD[114].

The initial state of HICs is sensitive to lots of interesting physics, including the nuclear structure and the saturation of the gluon distribution at small Bjorken-x. However, not all information of the initial state can be reconstructed based on final-state experimental observables because entropy is produced during the QGP evolution, which implies the loss of information. Nevertheless, recovering even part of the initial state information can still help to understand nuclear matter properties or collision patterns within the dynamic evolution. For example, attempts using Bayesian inference have been made to probe the nuclear structure[115] or neutron skin[116]. One essential initial state parameter, although a simple one, is the impact parameter. It governs the event geometry and volume estimation. Unlike early attempts using ML algorithms such as multilayer perceptrons or support vector machines to determine the impact parameter from conventional observables[117–119], a recent PointNet-based development[24,120] took advantage of the inherent point cloud structure of the detector output in HICs. Such an advance allows a real-time, end-to-end, event-wise impact parameter determination for the next generation of low-energy HICs to be performed at the compressed baryonic matter detector of the Facility for Antiproton and Ion Research.

Conclusions and outlook

Exploring QCD physics as inverse problems presents numerous challenges that drive innovative approaches to decode complex phenomena. Machine learning techniques, especially Bayesian inference and deep learning, have substantially advanced our ability to tackle these problems. Physics-driven designs ensure that solutions align with physical realities, which enhances precision and relevance. Embedding symmetries, physics principles and physical equations into deep models opens new avenues for reliable deep learning. Physics-driven learning has reformed first-principle calculations of lattice QCD and influenced investigations in hadron physics, neutron stars and heavy-ion collisions. Although this Perspective does not discuss advancements in holographic QCD, jet physics and electron ion collider owing to space constraints[121], it is worth noting here that anti-de Sitter space/conformal field theory correspondence has intrinsic connections with deep learning[122]. This relationship offers new ways to develop holographic theories using neural networks within physics-driven frameworks[123,124].

Several challenges with applying physics-driven learning to solve inverse problems remain:

- Demanding high-quality data: Solving inverse problems by linking experimental data to physical theories is particularly difficult because of the vast, unlabelled data generated by large-scale infrastructures (for example, collider detectors, astronomical observatories and high-performance computing simulations). Extracting valuable information from these data is non-trivial, especially owing to their noisy or incomplete nature. Moreover, the inherent difficulties in obtaining reliable observations during the measurement stages, which can introduce errors that propagate throughout the analysis, complicate the evaluation of uncertainties in these models.

- Reliability of physics models: Physics-driven learning cannot exceed the underlying physical priors they are built on. These priors are often based on approximations or simplified assumptions about the real world, and discrepancies between the model and actual physical phenomena will lead to biases.
- Reusability of physics-driven designs: Although physics-driven learning is often tailored to specific problems, one challenge is the ability to generalize and reuse these models across different domains. A design that offers some generality while retaining the advantages provided by following specific physical principles requires further efforts.

To address these challenges, research needs to focus on gathering high-quality data from experiments, on enhancing the understanding of physics through observations, and on incorporating more general physics-driven designs into deep learning models. Overall, this approach can create a positive feedback loop wherein the development of physics-driven learning methods and the solution of inverse problems reinforce each other, leading to continuous advancements in both fields. Despite these challenges, there are many future opportunities. Physics-driven learning can reduce computational costs and data requirements by narrowing the parameter space, which makes it a powerful tool for advancing applications. The integration of ML with physics has the potential to revolutionize the understanding of complex systems[125] beyond QCD, including applications in climate science to better predict weather patterns[126], in biophysics to understand active matter behaviours[127], and in astrophysics to understand cosmic phenomena[128].

Published online: 06 January 2025

References

1. Gross, F. et al. 50 years of quantum chromodynamics. *Eur. Phys. J. C* **83**, 1125 (2023).
2. Gattringer, C & Lang, C. B. *Quantum Chromodynamics on the Lattice* Vol. 788 (Springer, 2010).
3. Baym, G. et al. From hadrons to quarks in neutron stars: a review. *Rept. Prog. Phys.* **81**, 056902 (2018).
4. Yagi, K., Hatsuda, T. & Miake, Y. *Quark-Gluon Plasma: From Big Bang to Little Bang* Vol. 23 (Cambridge Univ. Press, 2005).
5. Tanaka, A., Tomiya, A. & Hashimoto, K. *Deep Learning and Physics* Vol. 1 (Springer, 2021).
6. Zhou, K., Wang, L., Pang, L.-G. & Shi, S. Exploring QCD matter in extreme conditions with machine learning. *Prog. Part. Nucl. Phys.* **135**, 104084 (2024).
7. Kaipio, J. & Somersalo, E. *Statistical and Computational Inverse Problems* Vol. 160 (Springer, 2006).
8. Asakawa, M., Hatsuda, T. & Nakahara, Y. Maximum entropy analysis of the spectral functions in lattice QCD. *Prog. Part. Nucl. Phys.* **46**, 459–508 (2001).
9. Yunes, N., Miller, M. C. & Yagi, K. Gravitational-wave and X-ray probes of the neutron star equation of state. *Nat. Rev. Phys.* **4**, 237–246 (2022).
10. Boehnlein, A. et al. Colloquium: machine learning in nuclear physics. *Rev. Mod. Phys.* **94**, 031003 (2022).
11. LeCun, Y., Bengio, Y. & Hinton, G. Deep learning. *Nature* **521**, 436–444 (2015).
12. Bishop, C. M. & Bishop, H. *Deep learning: Foundations and Concepts* (Springer Nature, 2023).
13. Raissi, M., Perdikaris, P. & Karniadakis, G. E. Physics informed deep learning (part I): data-driven solutions of nonlinear partial differential equations. *J. Comput. Phys.* **378**, 686–707 (2019).
14. Thuerey, N. et al. Physics-based deep learning. Preprint at https://doi.org/10.48550/arXiv.2109.05237 (2021).
15. Carleo, G. et al. Machine learning and the physical sciences. *Rev. Mod. Phys.* **91**, 045002 (2019).
16. LeCun, Y., Chopra, S., Hadsell, R., Ranzato, M. & Huang, F. in *Predicting Structured Data* 191–246 (MIT Press, 2007).
17. Jaynes, E. T. Information theory and statistical mechanics. *Phys. Rev.* **106**, 620–630 (1957).
18. Sohl-Dickstein, J., Weiss, E. A., Maheswaranathan, N. & Ganguli, S. Deep unsupervised learning using nonequilibrium thermodynamics. In *ICML'15: Proc. 32nd International Conference on International Conference on Machine Learning* (ICML, 2015).
19. Müller, B, Reinhardt, J & Strickland, M. T. *Neural Networks: An Introduction* (Springer, 1995).
20. Murphy, K. P. *Machine Learning: A Probabilistic Perspective* (MIT Press, 2012).
21. Jin, K. H., McCann, M. T., Froustey, E. & Unser, M. Deep convolutional neural network for inverse problems in imaging. *IEEE Trans. Image Process.* **26**, 4509–4522 (2017).
22. Amin, G. R. & Emrouznejad, A. Inverse forecasting: a new approach for predictive modeling. *Comput. Ind. Eng.* **53**, 491–498 (2007).
23. Behrmann, J., Grathwohl, W., Chen, R. T., Duvenaud, D. & Jacobsen, J.-H. Invertible residual networks. In *International Conference on Machine Learning* 573–582 (PMLR, 2019).
24. Qi, C. R., Su, H., Mo, K. & Guibas, L. J. PointNet: deep learning on point sets for 3D classification and segmentation. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition* 652–660 (IEEE, 2017).
25. Tomiya, A. & Nagai, Y. Equivariant transformer is all you need. *Proc. Sci.* **453**, 001 (2024).
26. Raissi, M., Perdikaris, P. & Karniadakis, G. E. Physics-informed neural networks: a deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J. Comput. Phys.* **378**, 686–707 (2019).
27. Mattheakis, M., Protopapas, P., Sondak, D., Di Giovanni, M. & Kaxiras, E. Physical symmetries embedded in neural networks. Preprint at https://doi.org/10.48550/arXiv.1904.08991 (2019).
28. Zhang, W., Tanida, J., Itoh, K. & Ichioka, Y. Shift-invariant pattern recognition neural network and its optical architecture. In *Proc. Annual Conference of the Japan Society of Applied Physics* 6p-M-14, 734 (JSAP, 1988).
29. Geiger, M. & Smidt, T. e3nn: Euclidean neural networks. Preprint at https://doi.org/10.48550/arXiv.2207.09453 (2022).
30. Shlomi, J., Battaglia, P. & Vlimant, J.-R. Graph neural networks in particle physics. *Mach. Learn. Sci. Technol.* **2**, 021001 (2020).
31. Kanwar, G. et al. Equivariant flow-based sampling for lattice gauge theory. *Phys. Rev. Lett.* **125**, 121601 (2020).
32. Favoni, M., Ipp, A., Müller, D. I. & Schuh, D. Lattice gauge equivariant convolutional neural networks. *Phys. Rev. Lett.* **128**, 032003 (2022).
33. Cranmer, K., Kanwar, G., Racanière, S., Rezende, D. J. & Shanahan, P. E. Advances in machine-learning-based sampling motivated by lattice quantum chromodynamics. *Nat. Rev. Phys.* **5**, 526–535 (2023).
34. Han, M.-Z., Jiang, J.-L., Tang, S.-P. & Fan, Y.-Z. Bayesian nonparametric inference of the neutron star equation of state via a neural network. *Astrophys. J.* **919**, 11 (2021).
35. Shi, S., Wang, L. & Zhou, K. Rethinking the ill-posedness of the spectral function reconstruction — why is it fundamentally hard and how artificial neural networks can help. *Comput. Phys. Commun.* **282**, 108547 (2023).
36. Karniadakis, G. E. et al. Physics-informed machine learning. *Nat. Rev. Phys.* **3**, 422–440 (2021).
37. Baydin, A. G., Pearlmutter, B. A., Radul, A. A. & Siskind, J. M. Automatic differentiation in machine learning: a survey. *J. Mach. Learn. Res.* **18**, 1–43 (2018).
38. Cranmer, K., Brehmer, J. & Louppe, G. The frontier of simulation-based inference. *Proc. Natl Acad. Sci. USA* **117**, 30055–30062 (2020).
39. Wang, H. et al. Scientific discovery in the age of artificial intelligence. *Nature* **620**, 47–60 (2023).
40. Jalali, B., Zhou, Y., Kadambi, A. & Roychowdhury, V. Physics-AI symbiosis. *Mach. Learn. Sci. Technol.* **3**, 041001 (2022).
41. Aarts, G. Introductory lectures on lattice QCD at nonzero baryon number. *J. Phys. Conf. Ser.* **706**, 022004 (2016).
42. Boyda, D. et al. Applications of machine learning to lattice quantum field theory. Preprint at https://doi.org/10.48550/arXiv.2202.05838 (2022).
43. Hasenfratz, P. & Niedermayer, F. Perfect lattice action for asymptotically free theories. *Nucl. Phys. B* **414**, 785–814 (1994).
44. DeGrand, T. A., Hasenfratz, A., Hasenfratz, P. & Niedermayer, F. The classically perfect fixed point action for SU(3) gauge theory. *Nucl. Phys. B* **454**, 587–614 (1995).
45. Shanahan, P. E., Trewartha, A. & Detmold, W. Machine learning action parameters in lattice quantum chromodynamics. *Phys. Rev. D* **97**, 094506 (2018).
46. Nagai, Y., Tanaka, A. & Tomiya, A. Self-learning Monte Carlo for non-Abelian gauge theory with dynamical fermions. *Phys. Rev. D* **107**, 054501 (2023).
47. Blücher, S., Kades, L., Pawlowski, J. M., Strodthoff, N. & Urban, J. M. Towards novel insights in lattice field theory with explainable machine learning. *Phys. Rev. D* **101**, 094507 (2020).
48. Holland, K., Ipp, A., Müller, D. I. & Wenger, U. Fixed point actions from convolutional neural networks. *Proc. Sci.* **453**, 038 (2024).
49. Holland, K., Ipp, A., Müller, D. I. & Wenger, U. Machine learning a fixed point action for SU(3) gauge theory with a gauge equivariant convolutional neural network. *Phys. Rev. D* **110**, 074502 (2024).
50. Zhou, K., Endrődi, G., Pang, L.-G. & Stöcker, H. Regressive and generative neural networks for scalar field theory. *Phys. Rev. D* **100**, 011501 (2019).
51. Pawlowski, J. M. & Urban, J. M. Reducing autocorrelation times in lattice simulations with generative adversarial networks. *Mach. Learn. Sci. Tech.* **1**, 045011 (2020).
52. Kanwar, G. Flow-based sampling for lattice field theories. In *40th International Symposium on Lattice Field Theory* (Proceedings of Science, 2024); https://doi.org/10.22323/1.453.0114.
53. Song, Y. et al. Score-based generative modeling through stochastic differential equations. In *NIPS '23: Proceedings of the 37th International Conference on Neural Information Processing Systems* 37799–37812 (ACM, 2020).
54. Parisi, G. & Wu, Y. S. Perturbation theory without gauge fixing. *Sci. Sin.* **24**, 483 (1980).
55. Wang, L., Aarts, G. & Zhou, K. Diffusion models as stochastic quantization in lattice field theory. *J. High Energy Phys.* **05**, 060 (2024).

56. Wang, L., Aarts, G. & Zhou, K. Generative diffusion models for lattice field theory. In *37th Conference on Neural Information Processing Systems* 21 (2023).
57. Zhu, Q., Aarts, G., Wang, W., Zhou, K. & Wang, L. Diffusion models for lattice gauge field simulations. In *38th conference on Neural Information Processing Systems* 14 (2024).
58. Hirono, Y., Tanaka, A. & Fukushima, K. Understanding diffusion models by Feynman's path integral. In *Proc. 41st International Conference on Machine Learning* (PMLR, 2024).
59. Cotler, J. & Rezchikov, S. Renormalizing diffusion models. Preprint at https://doi.org/10.48550/arXiv.2308.12355 (2023).
60. Müller, T., McWilliams, B., Rousselle, F., Gross, M. & Novák, J. Neural importance sampling. *ACM Trans. Graph.* **38**, 1–19 (2019).
61. Ron, D., Swendsen, R. H. & Brandt, A. Inverse Monte Carlo renormalization group transformations for critical phenomena. *Phys. Rev. Lett.* **89**, 275701 (2002).
62. Bachtis, D., Aarts, G., Di Renzo, F. & Lucini, B. Inverse renormalization group in quantum field theory. *Phys. Rev. Lett.* **128**, 081603 (2022).
63. Bachtis, D. Inverse renormalization group of spin glasses. *Phys. Rev. B* **110**, L140202 (2023).
64. Lehner, C. & Wettig, T. Gauge-equivariant neural networks as preconditioners in lattice QCD. *Phys. Rev. D* **108**, 034503 (2023).
65. Aronsson, J., Müller, D. I. & Schuh, D. Geometrical aspects of lattice gauge equivariant convolutional neural networks. Preprint https://doi.org/10.48550/arXiv.2303.11448 (2023).
66. Lehner, C. & Wettig, T. Gauge-equivariant pooling layers for preconditioners in lattice QCD. *Phys. Rev. D* **110**, 034517 (2023).
67. Cohen, T. S. & Welling, M. Group equivariant convolutional networks. In *Proceedings of The 33rd International Conference on Machine Learning* Vol. 48, 2990–2999 (PMLR, 2016).
68. Cohen, T. S., Weiler, M., Kicanaoglu, B. & Welling, M. Gauge equivariant convolutional networks and the icosahedral CNN. *Proc. 36th International Conference on Machine Learning* Vol. 97, 1321–1330 (PMLR, 2019).
69. Carrasquilla, J. & Melko, R. G. Machine learning phases of matter. *Nat. Phys.* **13**, 431–434 (2017).
70. Wetzel, S. J. & Scherzer, M. Machine learning of explicit order parameters: from the Ising model to SU(2) lattice gauge theory. *Phys. Rev. B* **96**, 184410 (2017).
71. Boyda, D. L. et al. Finding the deconfinement temperature in lattice Yang-Mills theories from outside the scaling window with machine learning. *Phys. Rev. D* **103**, 014509 (2021).
72. Lee, J. et al. Deep neural networks as Gaussian processes. Preprint at https://doi.org/10.48550/arXiv.1711.00165 (2018).
73. Halverson, J., Maiti, A. & Stoner, K. Neural networks and quantum field theory. *Mach. Learn. Sci. Tech.* **2**, 035002 (2021).
74. Bachtis, D., Aarts, G. & Lucini, B. Quantum field-theoretic machine learning. *Phys. Rev. D* **103**, 074510 (2021).
75. Aarts, G., Lucini, B. & Park, C. Scalar field restricted Boltzmann machine as an ultraviolet regulator. *Phys. Rev. D* **109**, 034521 (2024).
76. Aarts, G., Lucini, B. & Park, C. Stochastic weight matrix dynamics during learning and Dyson Brownian motion. Preprint at https://doi.org/10.48550/arXiv.2407.16427 (2024).
77. Rothkopf, A. Heavy quarkonium in extreme conditions. *Phys. Rept.* **858**, 1–117 (2020).
78. Wang, L., Shi, S. & Zhou, K. Reconstructing spectral functions via automatic differentiation. *Phys. Rev. D* **106**, L051502 (2022).
79. Guo, F.-K. et al. Hadronic molecules. *Rev. Mod. Phys.* **90**, 015004 (2018).
80. Sombillo, D. L. B., Ikeda, Y., Sato, T. & Hosaka, A. Classifying the pole of an amplitude using a deep neural network. *Phys. Rev. D* **102**, 016024 (2020).
81. Sombillo, D. L. B., Ikeda, Y., Sato, T. & Hosaka, A. Model independent analysis of coupled-channel scattering: a deep learning approach. *Phys. Rev. D* **104**, 036001 (2021).
82. Albaladejo, M. et al. Novel approaches in hadron spectroscopy. *Prog. Part. Nucl. Phys.* **127**, 103981 (2022).
83. Ng, L. et al. Deep learning exotic hadrons. *Phys. Rev. D* **105**, L091501 (2022).
84. Keeble, J. W. T. & Rios, A. Machine learning the deuteron. *Phys. Lett. B* **809**, 135743 (2020).
85. Adams, C., Carleo, G., Lovato, A. & Rocco, N. Variational Monte Carlo calculations of $A \leq 4$ nuclei with an artificial neural-network correlator ansatz. *Phys. Rev. Lett.* **127**, 022502 (2021).
86. Ishii, N., Aoki, S. & Hatsuda, T. The nuclear force from lattice QCD. *Phys. Rev. Lett.* **99**, 022001 (2007).
87. Aoki, S. et al. Lattice QCD approach to nuclear physics. *Prog. Theor. Exp. Phys.* **2012**, 01A105 (2012).
88. Aoki, S. & Doi, T. in *Handbook of Nuclear Physics* 1–31 (Springer, 2023).
89. Collaboration, A. et al. Unveiling the strong interaction among hadrons at the LHC. *Nature* **588**, 232–238 (2020).
90. Lyu, Y. et al. Doubly charmed tetraquark T_{cc}^{+} in (2+1)-flavor QCD near physical point. *Proc. Sci.* **453**, 077 (2024).
91. Shi, S., Zhou, K., Zhao, J., Mukherjee, S. & Zhuang, P. Heavy quark potential in the quark-gluon plasma: deep neural network meets lattice quantum chromodynamics. *Phys. Rev. D* **105**, 014017 (2022).
92. Wang, L., Doi, T., Hatsuda, T. & Lyu, Y. Building hadron potentials from lattice QCD with deep neural networks. Preprint at https://doi.org/10.48550/arXiv.2410.03082 (2024).
93. Fukushima, K., Mohanty, B. & Xu, N. Little-bang and femto-nova in nucleus-nucleus collisions. *AAPPS Bull.* **31**, 1 (2021).
94. Steiner, A. W., Lattimer, J. M. & Brown, E. F. The neutron star mass-radius relation and the equation of state of dense matter. *Astrophys. J. Lett.* **765**, L5 (2013).
95. Ozel, F. et al. The dense matter equation of state from neutron star radius and mass measurements. *Astrophys. J.* **820**, 28 (2016).
96. Brandes, L., Weise, W. & Kaiser, N. Inference of the sound speed and related properties of neutron stars. *Phys. Rev. D* **107**, 014011 (2023).
97. Fujimoto, Y., Fukushima, K. & Murase, K. Mapping neutron star data to the equation of state using the deep neural network. *Phys. Rev. D* **101**, 054016 (2020).
98. Carvalho, V., Ferreira, M., Malik, T. & Providência, C. Decoding neutron star observations: revealing composition through Bayesian neural networks. *Phys. Rev. D* **108**, 043031 (2023).
99. Carvalho, V., Ferreira, M. & Providência, C. From NS observations to nuclear matter properties: a machine learning approach. *Phys. Rev. D* **109**, 123038 (2024).
100. Soma, S., Wang, L., Shi, S., Stöcker, H. & Zhou, K. Neural network reconstruction of the dense matter equation of state from neutron star observables. *J. Cosmol. Astropart. Phys.* **08**, 071 (2022).
101. Soma, S., Wang, L., Shi, S., Stöcker, H. & Zhou, K. Reconstructing the neutron star equation of state from observational data via automatic differentiation. *Phys. Rev. D* **107**, 083028 (2023).
102. Fujimoto, Y., Fukushima, K. & Murase, K. Extensive studies of the neutron star equation of state from the deep learning inference with the observational data augmentation. *J. High Energy Phys.* **03**, 273 (2021).
103. Bass, S. A., Bernhard, J. E. & Moreland, J. S. Determination of quark-gluon-plasma parameters from a global Bayesian analysis. *Nucl. Phys. A* **967**, 67–73 (2017).
104. Pang, L.-G. et al. An equation-of-state-meter of quantum chromodynamics transition from deep learning. *Nature Commun.* **9**, 210 (2018).
105. Du, Y.-L. et al. Identifying the nature of the QCD transition in heavy-ion collisions with deep learning. *Nucl. Phys. A* **1005**, 121891 (2021).
106. Steinheimer, J. et al. A machine learning study on spinodal clumping in heavy ion collisions. *Nucl. Phys. A* **1005**, 121867 (2021).
107. Jiang, L., Wang, L. & Zhou, K. Deep learning stochastic processes with QCD phase transition. *Phys. Rev. D* **103**, 116023 (2021).
108. Omana Kuttan, M., Zhou, K., Steinheimer, J., Redelbach, A. & Stoecker, H. An equation-of-state-meter for CBM using PointNet. *J. High Energy Phys.* **21**, 184 (2020).
109. Pratt, S., Sangaline, E., Sorensen, P. & Wang, H. Constraining the eq. of state of super-hadronic matter from heavy-ion collisions. *Phys. Rev. Lett.* **114**, 202301 (2015).
110. Omana Kuttan, M., Steinheimer, J., Zhou, K. & Stoecker, H. QCD equation of state of dense nuclear matter from a Bayesian analysis of heavy-ion collision data. *Phys. Rev. Lett.* **131**, 202303 (2023).
111. Bernhard, J. E., Moreland, J. S. & Bass, S. A. Bayesian estimation of the specific shear and bulk viscosity of quark–gluon plasma. *Nat. Phys.* **15**, 1113–1117 (2019).
112. Everett, D. et al. Phenomenological constraints on the transport properties of QCD matter with data-driven model averaging. *Phys. Rev. Lett.* **126**, 242301 (2021).
113. Nijs, G., van der Schee, W., Gürsoy, U. & Snellings, R. Transverse momentum differential global analysis of heavy-ion collisions. *Phys. Rev. Lett.* **126**, 202301 (2021).
114. Li, F.-P., Lü, H.-L., Pang, L.-G. & Qin, G.-Y. Deep-learning quasi-particle masses from QCD equation of state. *Phys. Lett. B* **844**, 138088 (2023).
115. Cheng, Y.-L., Shi, S., Ma, Y.-G., Stöcker, H. & Zhou, K. Examination of nucleon distribution with Bayesian imaging for isobar collisions. *Phys. Rev. C* **107**, 064909 (2023).
116. Giacalone, G., Nijs, G. & van der Schee, W. Determination of the neutron skin of Pb208 from ultrarelativistic nuclear collisions. *Phys. Rev. Lett.* **131**, 202302 (2023).
117. David, C., Freslier, M. & Aichelin, J. Impact parameter determination for heavy-ion collisions by use of a neural network. *Phys. Rev. C* **51**, 1453–1459 (1995).
118. Bass, S. A., Bischoff, A., Maruhn, J. A., Stoecker, H. & Greiner, W. Neural networks for impact parameter determination. *Phys. Rev. C* **53**, 2358–2363 (1996).
119. De Sanctis, J. et al. Classification of the impact parameter in nucleus-nucleus collisions by a support vector machine method. *J. Phys. G* **36**, 015101 (2009).
120. Omana Kuttan, M., Steinheimer, J., Zhou, K., Redelbach, A. & Stoecker, H. A fast centrality-meter for heavy-ion collisions at the CBM experiment. *Phys. Lett. B* **811**, 135872 (2020).
121. Feickert, M. & Nachman, B. A living review of machine learning for particle physics. Preprint at https://doi.org/10.48550/arXiv.2102.02770 (2021).
122. Hashimoto, K., Sugishita, S., Tanaka, A. & Tomiya, A. Deep learning and the AdS/CFT correspondence. *Phys. Rev. D* **98**, 046019 (2018).
123. Hashimoto, K., Sugishita, S., Tanaka, A. & Tomiya, A. Deep learning and holographic QCD. *Phys. Rev. D* **98**, 106014 (2018).
124. Cai, R.-G., He, S., Li, L. & Zeng, H.-A. QCD phase diagram at finite magnetic field and chemical potential: a holographic approach using machine learning. Preprint at https://doi.org/10.48550/arXiv.2406.12772 (2024).
125. Kadambi, A., de Melo, C., Hsieh, C.-J., Srivastava, M. & Soatto, S. Incorporating physics into data-driven computer vision. *Nat. Mach. Intell.* **5**, 572–580 (2023).
126. Reichstein, M. et al. Deep learning and process understanding for data-driven earth system science. *Nature* **566**, 195–204 (2019).
127. Cichos, F., Gustavsson, K., Mehlig, B. & Volpe, G. Machine learning for active matter. *Nat. Mach. Intell.* **2**, 94–103 (2020).
128. Huerta, E. A. et al. Enabling real-time multi-messenger astrophysics discoveries with deep learning. *Nat. Rev. Phys.* **1**, 600–608 (2019).

Acknowledgements
The authors thank L. Brandes, Y. Fujimoto, S. Kamata, D. I. Müller, K. Murase, A. Tanaka and A. Tomiya for the helpful discussions and all collaborators for their contributions. The authors also thank the DEEP-IN working group at RIKEN-iTHEMS, the European Centre for Theoretical Studies in Nuclear Physics and Related Areas (ECT*), and the ExtreMe Matter Institute (EMMI) for their support in the preparation of this paper. G.A. is supported by STFC Consolidated

Grant ST/T000813/1. K.F. is supported by Japan Society for the Promotion of Science (JSPS) KAKENHI Grant numbers 22H01216 and 22H05118. T.H. is supported by the Japan Science and Technology Agency (JST) as part of Adopting Sustainable Partnerships for Innovative Research Ecosystem (ASPIRE), grant number JPMJAP2318. S.S. acknowledges Tsinghua University under grant number 53330500923. L.W. thanks the National Natural Science Foundation of China (number 12147101) for supporting his visit to Fudan University. K.Z. is supported by the CUHK-Shenzhen University Development Fund under grant numbers UDF01003041 and UDF03003041, and Shenzhen Peacock fund under grant number 2023TC0179.

Author contributions
The authors contributed equally to all aspects of the article.

Competing interests
The authors declare no competing interests.

Additional information
Peer review information *Nature Reviews Physics* thanks Long-Gang Pang and Danilo Jimenez Rezende for their contribution to the peer review of this work.