Nigerian Journal of Physics (NJP)
ISSN online: 3027-0936
ISSN print: 1595-0611
DOI: https://doi.org/10.62292/njp.v35(s).2026.781
Volume 35(S), 2026

Artificial Intelligence in Physics Research: Foundations, Applications and Challenges for Autonomous Scientific Discovery

[*1]Chijioke Onyia, [1]Chukwujekwu N. Ofodum, [1]Emeka M. Madike, [1]James Chibuike Ike, [1]Augustina Obiageli Chukwudi, [1]Onyinyechi Jennifer Ugwueze, [1]Chidi Gloria Ugwueye

[1]National Space Research and Development Agency (NASRDA), Abuja, Nigeria

*Corresponding Author's Email: onyiacj88@gmail.com

Keywords:
Artificial Intelligence, Machine Learning, Neural Networks, Particle Physics, Gravitational-Wave Astronomy, Materials Discover And Physics-Informed Neural Networks

Submitted: 11/11/2026

Accepted: 21/11/2026

Published: 24/11/2026

ABSTRACT

Artificial intelligence (AI) is increasingly used across physics research, but claims that it now produces autonomous scientific discovery require careful scrutiny. This narrative review aims to synthesize evidence published between January 2014 and August 2026, together with earlier foundational works, on how AI is used in physics and to evaluate the extent to which its contributions amount to validated discovery rather than prediction. Literature was identified through Scopus, Web of Science, Google Scholar, arXiv and INSPIRE-HEP and was selected for relevance to four areas: the statistical-physics foundations of machine learning; particle physics and gravitational-wave astronomy; materials discovery and quantum chemistry; and physics-informed neural networks (PINNs). The evidence shows that machine learning is now embedded in collider event reconstruction and jet tagging, has accelerated gravitational-wave detection and Bayesian parameter estimation by several orders of magnitude in published demonstrations and has greatly expanded computational searches for stable crystals; the Graph Networks for Materials Exploration (GNoME) model, for example, predicted 381,000 new stable inorganic crystals, of which 736 had been independently realized by external experimental groups at the time of publication. However, most reported advances remain computational predictions or benchmark demonstrations, several of the most ambitious claims are still preprints and the interpretability, reproducibility and experimental validation of AI outputs vary widely across subfields. The review concludes that AI is a powerful accelerator of physics research but that autonomous theory discovery is not yet routine and that standardized validation, interpretable models and a clear distinction between prediction and confirmed discovery are needed before such claims can be generalized.

INTRODUCTION

The relationship between physics and artificial intelligence was long largely one-directional such that physics supplied concepts that AI research adopted. This is changing, AI methods are now not only applied to physics problems but are also producing results that are increasingly presented as independent scientific contributions. When the 2024 Nobel Prize in Physics was awarded to John Hopfield and Geoffrey Hinton for foundational work on artificial neural networks, it was a public acknowledgment that the exchange between these fields has reshaped both of them (Royal Swedish Academy of Sciences, 2024).

Hopfield's associative memory network, introduced in 1982, borrowed directly from spin-system physics to store and retrieve patterns in data (Hopfield, 1982). Hinton extended this approach with the Boltzmann machine, using the energy dynamics of statistical physics to show that neural networks could learn internal representations from examples (Ackley et al., 1985; Royal Swedish Academy of Sciences, 2024). Both ideas originated in statistical mechanics and both became foundations of the deep-learning methods that have since spread across most areas of science.

The pace of adoption is evident in the literature. The HEP ML Living Review, which tracks machine-learning papers in high-energy physics, now records hundreds of

publications per year on topics including jet tagging, event reconstruction, anomaly detection and simulation acceleration (Krause *et al*., 2026). In materials science, AI frameworks have predicted hundreds of thousands of candidate stable crystal structures, a small fraction of which correspond to experimentally realized compounds (Merchant et al., 2023). In gravitational-wave astronomy, machine learning is used for signal detection, parameter estimation and noise characterization (Cuoco et al., 2025).

Several reviews have catalogued these developments within individual subfields, for example in particle physics (Krause et al., 2026), gravitational-wave research (Cuoco et al., 2025), materials informatics (Lookman et al., 2026) and PINNs (Cuomo et al., 2022). Less attention has been given to comparing, across subfields, how reliable, interpretable, reproducible and experimentally validated AI-derived results are, or to distinguishing clearly between prediction, experimental confirmation and explanatory understanding. This distinction matters because accounts of "AI discovery" frequently treat computational predictions as though they were confirmed findings (Cheetham & Seshadri, 2024).

This review covers literature published through August 2026 and is organized as follows. Section 2 describes the review methodology. Section 3 examines the statistical-physics foundations of machine learning. Section 4 discusses applications in particle physics and gravitational-wave astronomy, Section 5 addresses materials discovery and quantum chemistry and Section 6 considers physics-informed neural networks and their quantum extensions. Section 7 presents a comparative synthesis of the evidence (Table 1), Section 8 examines whether these tools produce new physical understanding or sophisticated pattern recognition and Section 9 concludes.

The aim of this review is to critically synthesize the evidence on AI-driven research in physics across four subfields and to evaluate, using explicit criteria of model reliability, interpretability, reproducibility and validation status, the extent to which AI contributions constitute validated scientific discovery rather than predictive performance.

MATERIALS AND METHODS

Review Methodology

Review design

This article is a narrative (non-systematic) review. It did not follow a registered systematic-review protocol such as PRISMA and no quantitative meta-analysis was performed. A narrative design was chosen because the topic spans subfields with different publication cultures, evaluation metrics and validation standards that are not amenable to pooled statistical analysis.

Sources and Search Strategy

Literature was identified through Scopus, the Web of Science Core Collection, Google Scholar, the arXiv preprint server and the INSPIRE-HEP database, supplemented by backward and forward citation tracking from key papers and by the HEP ML Living Review (Krause et al., 2026). Searches combined a method term ("artificial intelligence", "machine learning", "deep learning", "neural network", "graph neural network" or "physics-informed neural network") with a domain term ("particle physics", "Large Hadron Collider", "jet tagging", "gravitational wave", "LIGO", "materials discovery", "crystal structure prediction", "interatomic potential", "Schrödinger equation", "partial differential equation" or "quantum"). For Section 8, method terms were also combined with "scientific discovery", "scientific understanding" and "interpretability".

Publication Period

The principal search window was January 2014 to August 2026. Earlier works were included only where they are foundational primary sources for concepts discussed in the review (e.g., Ackley et al., 1985; Hopfield, 1982).

Inclusion and Exclusion Criteria

Works were included if they (a) applied, developed or analysed AI or machine-learning methods for a problem in physics or a closely related physical science; (b) reported a quantitative result, a methodological advance or a substantive conceptual argument; and (c) were written in English. Peer-reviewed journal articles, conference proceedings, book chapters and preprints were eligible. Works that used AI only incidentally, duplicate versions of the same study (the peer-reviewed version was retained where available) and opinion pieces without supporting evidence were excluded.

Screening Procedure

Titles and abstracts were screened for relevance by the authors, followed by full-text reading of retained items; uncertainties about inclusion were resolved by discussion among the authors. Priority was given to primary sources reporting original results. Reviews were used to map each subfield and to locate primary studies.

Treatment of Preprints and News Reports

Preprints were included where they represent recent developments not yet available in peer-reviewed form; they are identified as preprints in the text and in Table 1 and their claims are treated as provisional. News reports, press releases and organizational webpages were used only to locate primary studies or to document institutional events (e.g., the Nobel Prize announcement) and company-reported figures, which are identified as

such. No central quantitative claim in this review rests on a news report alone.

Approach to Synthesis

The evidence was synthesized thematically and comparatively. For each application, the physics task, model class, dataset or benchmark, validation status, main limitation and original reference were extracted. These are summarized in Table 1 and discussed in Section 7 with reference to four comparative questions: How reliable is the model's performance? How interpretable are its outputs? How reproducible is the result? Has the result been validated experimentally, or only computationally?

Foundations: Physics Meets Machine Learning

From Spin Systems to Neural Networks

In an important sense, modern machine learning originated in physicists' work on magnetic systems. Hopfield's 1982 network recast associative memory as a problem in disordered spin systems, in which each neuron behaves like a binary spin variable and the network has an energy function whose minima are the stored patterns (Barra et al., 2026; Hopfield, 1982). When neurons update their states to lower the energy, the process is formally analogous to a physical system relaxing toward equilibrium. Statistical mechanics therefore supplied both the mathematical formalism and the intuition for understanding these networks.

Hinton extended the physical analogy further. His Boltzmann machine added stochastic dynamics governed by a Boltzmann distribution, which allowed the network to explore its energy landscape instead of becoming trapped in local minima (Ackley et al., 1985). Training the machine meant tuning connection weights until the network's equilibrium distribution matched the training data; this procedure relies on thermodynamic equilibrium and detailed balance. The Nobel Committee stated that both laureates used tools from statistical physics, the science of systems composed of many similar parts, to develop methods that underpin present-day machine learning (Royal Swedish Academy of Sciences, 2024).

The connections have since deepened. The renormalization group, one of the central tools of theoretical physics, has been shown to map exactly onto a form of hierarchical feature extraction performed by deep networks built from restricted Boltzmann machines (Mehta & Schwab, 2014). Tensor network methods built for quantum many-body physics now appear in machine learning architectures. Information-theoretic concepts such as entropy, mutual information and the information bottleneck serve both as analytical tools and as training objectives in deep learning. The 2024 Nobel Prize therefore recognized not a one-way transfer of ideas from physics to computer science, but a long-standing mutual influence between the two fields.

The Significance of The 2024 Nobel Prize in Physics

Ellen Moons, who chaired the Nobel Committee for Physics, noted that neural networks have advanced work in particle physics, materials science and astrophysics (Physics World, 2024). The prize did not honour Hopfield networks or Boltzmann machines as practical tools, since both are modest by current deep-learning standards; rather, it recognized the principle that physical laws can serve as the design basis for learning systems.

Hinton, speaking at the prize announcement, said that AI will reshape society on a scale comparable to the industrial revolution, adding the caveat that humanity has no experience of things that might be more intelligent than humans (Hao, 2025). One interpretation of the prize is that the boundary between physics and computer science has become increasingly permeable: physics supplies principles that make learning possible and learning systems are increasingly used to conduct physics research.

RESULTS AND DISCUSSION

AI in Particle Physics and Gravitational-Wave Astronomy

Machine learning at the Large Hadron Collider

Particle physicists were early adopters and neural networks were used in event classification as early as the late 1980s (Boge & de Regt, 2025). What has changed is the scale. Faster hardware and deeper architectures have greatly expanded the range of tasks to which machine learning is applied in the field (Boge & de Regt, 2025). The *HEP ML Living Review* now records dozens of new papers every month covering jet tagging, track reconstruction, anomaly detection, simulation and direct physics analysis (Krause et al., 2026).

Applications now span the entire analysis chain. Graph neural networks have become a widely used approach to track reconstruction because detector hits map naturally onto graph nodes, with edges representing possible particle trajectories (Ju et al., 2020). Transformer architectures, borrowed from natural language processing, have been adapted for jet tagging, that is, identifying the particle that produced a given spray of particles in a high-energy collision (Krause et al., 2026). More recent work has reported neural scaling laws for jet generation and tested multimodal approaches that fuse different data representations (Krause et al., 2026). Since most of these models are trained on simulated events, their performance on real collision data depends on how well the simulation is calibrated and the associated systematic uncertainties must be quantified separately for each application.

Two recent studies illustrate claims that go beyond prediction. Alexander et al. (2026) presented a neuro-symbolic system designed to derive particle physics theories directly from experimental data without a human theorist in the loop. Because this work is currently an arXiv preprint that has not been peer reviewed, it demonstrates a possibility rather than an established capability. Yu et al. (2025), writing in the *Proceedings of the National Academy of Sciences*, built a physics-tailored machine learning model that inferred force laws in a laboratory dusty plasma, describing non-reciprocal interparticle forces with an accuracy exceeding 99% and challenging long-standing assumptions about how these forces behave. Importantly, the model's inferences were validated against laboratory experiments.

Machine Learning in Gravitational-Wave Astronomy

Gravitational-wave astronomy is in some respects an ideal test case for machine learning. The signals, which are spacetime ripples from merging black holes, neutron stars and other energetic astrophysical events, are buried in detector noise that drifts, contains transient glitches and departs from Gaussian statistics in ways that are difficult to model (Cuoco et al., 2025). Standard matched-filter searches require large template banks and substantial computing resources, which is precisely the kind of bottleneck that data-driven methods aim to relieve.

At LIGO and its partner detectors, machine learning is now applied to detection, characterization and interpretation (Cuoco et al., 2025). Deep residual networks have reached sensitive distances that, in a specific search setting, exceed those of traditional techniques at a small fraction of the computational cost of matched filtering, making them candidates for rapid online searches that could speed up alerts for electromagnetic follow-up (Nousi et al., 2023). For Bayesian parameter estimation, conventional sampling-based analyses have typically taken between hours and days per event; a conditional variational autoencoder generated posterior samples for binary black hole signals around six orders of magnitude faster (Gabbard et al., 2022) and neural posterior estimation produced results within seconds that agree closely with standard inference for events in the first LIGO–Virgo transient catalogue (Dax et al., 2021). These speed-ups were demonstrated for specific source classes and prior ranges and the networks must be retrained or re-validated when detector noise or the signal model changes. The Gravity Spy project takes a hybrid approach, pairing citizen scientists with machine learning classifiers to classify detector glitches and reduce false alarm rates (Cuoco et al., 2025).

A University of California, Riverside team has developed an unsupervised method for mining LIGO's auxiliary channel data to improve the detector's self-diagnostics (ScienceDaily, 2025). Data from the first part of the fourth LIGO–Virgo–KAGRA observing run have also been used for direct searches for ultralight dark matter (LIGO Scientific Collaboration, Virgo Collaboration, & KAGRA Collaboration, 2026). Although that search relies on established signal-processing methods rather than machine learning, it illustrates the widening range of analyses that detector data must support and, consequently, the growing demand for scalable data-driven tools.

AI for Materials Discovery and Quantum Chemistry

Gnome and The Acceleration of Materials Knowledge

GNoME is among the most widely cited examples of AI applied to materials science. Google DeepMind's Graph Networks for Materials Exploration (a message-passing graph neural network trained on density functional theory data) identified 2.2 million crystal structures below the previously known convex hull, of which 381,000 were reported as new stable materials on the updated convex hull (Merchant et al., 2023). These figures refer to thermodynamic stability computed with density functional theory, not to synthesized materials. The same study reported that 736 of its stable structures had been independently realized experimentally by external groups; these are compounds made independently of the model, which the model also identified, rather than materials synthesized as a result of its predictions (Merchant et al., 2023). The final GNoME models predicted energies with an error of about 11 meV/atom on relaxed structures and achieved a precision of stable predictions (hit rate) above 80% when structural information was available (Merchant et al., 2023).

The developers also reported candidate materials of technological interest, including about 52,000 layered compounds similar to graphene and 528 potential lithium-ion conductors (Merchant & Cubuk, 2023). These are company-reported computational candidates and their stability under realistic conditions and their functional properties remain to be confirmed experimentally.

The experimental significance of these numbers has been debated. Cheetham and Seshadri (2024) examined a sample of the new GNoME compounds and found scant evidence of compounds that are simultaneously novel, credible and useful. In a companion study, an autonomous laboratory (A-Lab) reported the robotic synthesis of computationally targeted compounds over 17 days of continuous operation (Szymanski et al., 2023); a subsequent re-analysis of the diffraction data concluded that the novelty of these products had not been demonstrated (Leeman et al., 2024) and the original authors later revised the article to clarify that "novel" meant new to the prediction platform rather than

new to science. These exchanges illustrate why computational stability, independent synthesis and confirmed novelty must be reported as distinct levels of evidence.

Lookman et al. (2026), writing in *Advanced Materials*, traced materials informatics from its roots in physics and information theory through its current phase of AI-driven autonomous discovery. A central development they identify is the self-driving laboratory, which couples AI-guided experimental design with robotic synthesis and characterization so that the design–make–test cycle can proceed with minimal human intervention. Separately, Zeni et al. (2025) published a generative model for inorganic materials design in *Nature* that goes beyond screening known chemical spaces; it generates new material structures conditioned on target properties and the authors synthesized one generated structure whose measured property was within 20% of the target value.

Deep Learning for Computational Chemistry

Machine learning is changing computational chemistry from the electronic structure up. Deep neural networks can now serve as wavefunction ansätze for the electronic Schrödinger equation: The FermiNet architecture (Pfau et al., 2020) reaches chemical accuracy for a range of small molecules and atoms through variational Monte Carlo. At a coarser level, machine-learned force fields serve as fast surrogates for ab initio calculations, allowing molecular dynamics at near-quantum accuracy over timescales and system sizes that direct quantum computation cannot reach.

AlphaFold, developed by Google DeepMind, was recognized when Demis Hassabis and John Jumper shared the 2024 Nobel Prize in Chemistry for protein structure prediction (Schlembach & Wrublewski, 2025). AlphaFold 3 extended prediction to complexes of proteins with nucleic acids, small molecules, ions and modified residues and reported substantially higher accuracy than specialized tools for protein–ligand and protein–nucleic acid interactions (Abramson et al., 2024). The AlphaFold Protein Structure Database provides open access to hundreds of millions of predicted structures and received a redesigned interface in 2025 (Bertoni et al., 2026). That success has motivated parallel efforts in materials science; universal machine learning interatomic potentials such as the UMA model family (Wood et al., 2025), trained on about half a billion atomic structures, aim to provide transferable force fields across molecules, materials and catalysts.

Physics-Informed Neural Networks and Hybrid Architectures

The PINN Paradigm

Physics-informed neural networks (PINNs), introduced by Raissi et al. (2019), take a different approach to the trade-off between data and physical models. Instead of learning from data alone, a PINN incorporates the governing partial differential equations of a physical system into the loss function, penalizing solutions that violate known physics. The approach has spread into fluid dynamics, materials modelling, quantum mechanics and climate science (Cuomo et al., 2022; Okpechi-Kanayochukwu et al., 2026).

PINNs can solve forward and inverse problems, perform well where data are scarce or costly and enforce conservation laws that a purely data-driven model might violate. However, they also have well-documented weaknesses. Multiscale problems, sharp gradients and high-dimensional systems remain difficult for PINNs and the automatic differentiation needed to compute physics residuals adds computational overhead, so that for well-posed forward problems conventional numerical solvers are often faster (Cuomo et al., 2022).

Quantum Extensions: QPINNs

A more recent line of work seeks to run physics-informed learning on quantum hardware. Quantum physics-informed neural networks (QPINNs) combine quantum circuits with the PINN framework, motivated by the hypothesis that the exponential representational capacity of quantum states could encode high-dimensional functions in superposition (Panichi et al., 2026). Panichi, Corli and Prati, writing in *Physical Review Applied* in January 2026, showed QPINNs solving multivariable partial differential equations, whereas previous versions had largely been limited to ordinary differential equations.

Berger et al. (2025) added trainable embeddings to QPINNs for nonlinear partial differential equations (PDEs) and found that the quantum version could match classical PINN accuracy with fewer parameters in some cases. Hybrid quantum–classical PINN designs have been tested on computational fluid dynamics in complex geometries (Sedykh et al., 2024) and wavelet-accelerated QPINNs have been proposed, in a preprint, to handle the multiscale features that challenge standard approaches (Gupta et al., 2025). These studies report proof-of-concept results on comparatively small problem instances and none of the quantum extensions has yet demonstrated a clear advantage over classical PINNs. They should therefore be regarded as methodological demonstrations rather than evidence of a quantum advantage.

Comparative Synthesis of the Evidence

Table 1 summarizes the principal applications discussed in this review by model, dataset or benchmark, validation status, main limitation and original reference. Three levels of validation status are distinguished: (i)

computational or benchmark demonstration, in which performance is assessed on simulated data or held-out test sets; (ii) validation on real observational or experimental data against established methods; and (iii) independent experimental confirmation of a new physical result. Studies available only as preprints are marked as such.

Table 1: Comparative Summary of Representative AI Applications in Physics

Application	Model	Dataset / benchmark	Validation status	Main limitation	Original reference
Track reconstruction (collider physics)	Graph neural network	Simulated tracking-detector and calorimeter data	(i)	Dependence on simulation; scaling to high-luminosity conditions	Ju et al. (2020) [preprint]
Jet tagging and jet generation	Transformers; generative models	Simulated and collision data	(i)–(ii)	Calibration and systematic uncertainties inherited from simulation	Krause et al. (2026) [preprint; living review]
Theory extraction from collider data	Neuro-symbolic system	Experimental particle physics data	(i)	Not yet peer reviewed; restricted theory space	Alexander et al. (2026) [preprint]
Force laws in dusty plasma	Physics-tailored neural network	3D particle trajectories from laboratory experiments	(iii)	Single experimental system; generality untested	Yu et al. (2025)
GW signal detection	Deep residual network (AResGW)	Injected signals in real O3a detector noise	(ii)	Advantage shown for a specific setting; retraining needed as noise changes	Nousi et al. (2023)
GW parameter estimation	Conditional variational autoencoder; neural posterior estimation	Simulated BBH signals; first LIGO–Virgo catalogue events	(ii)	Restricted source classes and priors	Gabbard et al. (2022); Dax et al. (2021)
Crystal stability prediction	GNN (GNoME)	Materials Project, OQMD and DFT calculations	(i); 736 structures independently realized	Computed stability is not synthesizability; novelty contested	Merchant et al. (2023); Cheetham & Seshadri (2024)
Autonomous synthesis	Robotic laboratory with ML-guided planning (A-Lab)	Computationally screened target compounds	(iii) claimed; contested	Phase identification and novelty disputed	Szymanski et al. (2023); Leeman et al. (2024)
Generative materials design	Diffusion model (MatterGen)	Inorganic crystal-structure databases	(i), with one synthesized structure	Limited experimental validation	Zeni et al. (2025)
Electronic structure	FermiNet wavefunction ansatz	Small atoms and molecules	(ii) against high-accuracy quantum chemistry	Computational cost grows steeply with system size	Pfau et al. (2020)
Biomolecular	AlphaFold 3	Protein Data Bank	(ii) against	Limited physical	Abramson et al.

structure		structures	experimental structures	interpretability	(2024)
Interatomic potentials	UMA model family	~500 million DFT-labelled structures	(i)–(ii)	Inherits DFT errors; reliability outside training domain	Wood et al. (2025)
Solving PDEs	PINN	Canonical forward and inverse PDE problems	(i)	Multiscale and stiff problems; training cost	Raissi et al. (2019); Cuomo et al. (2022)
Quantum-enhanced PDE solving	Variational quantum circuits (QPINNs)	Small PDE instances	(i)	No demonstrated quantum advantage	Panichi et al. (2026); Berger et al. (2025); Gupta et al. (2025) [preprint]

Note. Validation status: (i) computational or benchmark demonstration; (ii) validation on real data against established methods; (iii) independent experimental confirmation. BBH = binary black hole; DFT = density functional theory; GNN = graph neural network; GW = gravitational wave; PDE = partial differential equation.

In particle physics and gravitational-wave astronomy, machine-learning models are mostly validated against established methods on simulated or real detector data; their reliability is well characterized for specific tasks, but their dependence on simulation and on training priors limits generalization. In materials science, the scale of computational prediction far exceeds the scale of experimental confirmation and the most widely publicized claims of discovery have been contested on grounds of novelty and characterization. PINNs and their quantum extensions remain largely at the stage of methodological demonstration. Interpretability is weakest for large discriminative and generative models and strongest for physics-constrained architectures, such as PINNs and the physics-tailored model of Yu et al. (2025), whose structure encodes known symmetries. Reproducibility is supported where code, data or model weights are released, as for the GNoME structures, the AResGW detection code and the UMA models (Merchant et al., 2023; Nousi et al., 2023; Wood et al., 2025), but independent replication remains uncommon. Across subfields, few of the studies reviewed reach the level of independent experimental confirmation of a new physical result and those that do concern narrowly defined systems.

Epistemological Challenges: Discovery Versus Pattern Recognition

These developments raise a fundamental question which is does AI produce new physical understanding, or does it perform highly efficient pattern recognition that is then described as discovery? Boge and de Regt (2025), in a chapter in *Philosophy of Science for Machine Learning*, argued that AI can make meaningful scientific discoveries without starting from an existing theory, while cautioning that this claim raises a number of open problems.

Some results are difficult to dismiss as pattern matching alone. Yu et al.'s (2025) dusty plasma work is a case in point where the physics-tailored neural network did not merely reproduce known force behaviours with high accuracy; it identified non-reciprocal interaction features that were unexpected and that challenged existing assumptions and these inferences were checked against laboratory measurements. GNoME is sometimes cited as a comparable case from materials science, but the evidence is weaker: its stable-crystal predictions are computational, the 736 experimentally realized structures were produced independently of the model and claims of novelty for autonomously synthesized compounds have been contested (Cheetham & Seshadri, 2024; Leeman et al., 2024). Predictive success at this scale is significant, but it does not by itself constitute experimental discovery.

Sceptics also raise a deeper objection. A neural network that accurately predicts material stability does not explain why a crystal structure is stable, at least not in the way a first-principles account does. A model can therefore be accurate without being transparent and predictive without being explanatory. Predictive power and explanatory depth have always been in tension in science, but deep learning widens the gap because the models are so large and so opaque. What AI-driven physics requires and largely does not yet have, are methods that go beyond finding patterns and distil them into compact, generalizable principles that physicists would recognize as understanding. Symbolic and neuro-symbolic approaches, such as the system described by Alexander et al. (2026), are one route toward this goal, although their outputs must be validated in the same way as any other theoretical proposal.

CONCLUSION

This narrative review has shown that AI is now embedded across physics research, from collider event

reconstruction and gravitational-wave inference to materials screening, electronic-structure calculation and the solution of differential equations. The established evidence is strongest for AI as an accelerator of existing workflows such as in published demonstrations, machine learning reduces the time required for gravitational-wave parameter estimation by several orders of magnitude, matches or exceeds traditional pipelines on specific detection tasks and greatly expands computational searches for stable materials. Evidence for AI as an independent source of validated physical discovery is much thinner; it rests on a small number of experimentally checked cases, such as the inference of dusty-plasma force laws, while the most prominent materials-discovery claims remain computational or contested and proposals for autonomous theory extraction are still preprints. Unresolved limitations include the opacity of large models, dependence on simulated training data, limited independent replication and the absence of common standards for reporting validation status. Realistic future directions therefore include interpretable and neuro-symbolic methods that translate learned patterns into testable physical principles; standardized validation and reproducibility criteria for AI-derived results, including explicit reporting of whether a result is a prediction, a benchmark demonstration or an experimentally confirmed finding; shared benchmark datasets that permit fair comparison of AI-driven and conventional methods across subfields; and clear guidelines for attributing credit in human–AI collaborative research. AI is already changing how physics is done, but whether it will change what physics understands will depend on how rigorously these standards are developed and applied.

REFERENCES

Abramson, J., Adler, J., Dunger, J., Evans, R., Green, T., Pritzel, A., Ronneberger, O., Willmore, L., Ballard, A. J., Bambrick, J., Bodenstein, S. W., Evans, D. A., Hung, C.-C., O'Neill, M., Reiman, D., Tunyasuvunakool, K., Wu, Z., Žemgulytė, A., Arvaniti, E., . . . Jumper, J. M. (2024). Accurate structure prediction of biomolecular interactions with AlphaFold 3. *Nature*, *630*(8016), 493–500. https://doi.org/10.1038/s41586-024-07487-w

Ackley, D. H., Hinton, G. E., & Sejnowski, T. J. (1985). A learning algorithm for Boltzmann machines. *Cognitive Science*, *9*(1), 147–169. https://doi.org/10.1207/s15516709cog0901_7

Alexander, S., Bradley, B., Gouskos, L., & Niu, C. (2026). *Autonomous discovery of particle physics theories from experimental data* (arXiv:2603.28935). arXiv. https://doi.org/10.48550/arXiv.2603.28935

Barra, A., Durante, F., Ladiana, A., & Solazzo, M. M. (2026). *Do Hopfield networks dream of stored patterns? A statistical-mechanical theory of dreaming in multidirectional associative memories* (arXiv:2605.13721). arXiv. https://doi.org/10.48550/arXiv.2605.13721

Berger, S., Hosters, N., & Möller, M. (2025). Trainable embedding quantum physics informed neural networks for solving nonlinear PDEs. *Scientific Reports*, *15*(1), Article 18823. https://doi.org/10.1038/s41598-025-02959-z

Bertoni, D., Tsenkov, M., Magana, P., Nair, S., Pidruchna, I., Querino Lima Afonso, M., Midlik, A., Paramval, U., Lawal, D., Tanweer, A., Last, M., Patel, R., Laydon, A., Lasecki, D., Dietrich, N., Tomlinson, H., Žídek, A., Green, T., Kovalevskiy, O., . . . Velankar, S. (2026). AlphaFold Protein Structure Database 2025: A redesigned interface and updated structural coverage. *Nucleic Acids Research*, *54*(D1), D358–D362. https://doi.org/10.1093/nar/gkaf1226

Boge, F. J., & de Regt, H. W. (2025). Machine learning discoveries and scientific understanding in particle physics: Problems and prospects. In J. M. Durán & G. Pozzi (Eds.), *Philosophy of science for machine learning: Core issues and new perspectives* (pp. 403–434). Springer. https://doi.org/10.1007/978-3-032-03083-2_19

Cheetham, A. K., & Seshadri, R. (2024). Artificial intelligence driving materials discovery? Perspective on the article: Scaling deep learning for materials discovery. *Chemistry of Materials*, *36*(8), 3490–3495. https://doi.org/10.1021/acs.chemmater.4c00643

Cuoco, E., Cavaglià, M., Heng, I. S., Keitel, D., & Messenger, C. (2025). Applications of machine learning in gravitational-wave research with current interferometric detectors. *Living Reviews in Relativity*, *28*(1), Article 2. https://doi.org/10.1007/s41114-024-00055-8

Cuomo, S., Di Cola, V. S., Giampaolo, F., Rozza, G., Raissi, M., & Piccialli, F. (2022). Scientific machine learning through physics-informed neural networks: Where we are and what's next. *Journal of Scientific Computing*, *92*(3), Article 88. https://doi.org/10.1007/s10915-022-01939-z

Dax, M., Green, S. R., Gair, J., Macke, J. H., Buonanno, A., & Schölkopf, B. (2021). Real-time gravitational wave science with neural posterior estimation. *Physical Review Letters*, *127*(24), Article 241103. https://doi.org/10.1103/PhysRevLett.127.241103

Gabbard, H., Messenger, C., Heng, I. S., Tonolini, F., & Murray-Smith, R. (2022). Bayesian parameter estimation using conditional variational autoencoders for gravitational-wave astronomy. *Nature Physics, 18*(1), 112–117. https://doi.org/10.1038/s41567-021-01425-7

Gupta, D., Pandey, H., & Behera, R. (2025). *Wavelet-accelerated physics-informed quantum neural network for multiscale partial differential equations* (arXiv:2512.08256). arXiv. https://doi.org/10.48550/arXiv.2512.08256

Hao, K. (2025). *Empire of AI: Inside the reckless race for total domination*. Random House.

Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences, 79*(8), 2554–2558. https://doi.org/10.1073/pnas.79.8.2554

Ju, X., Farrell, S., Calafiura, P., Murnane, D., Prabhat, Gray, L., Klijnsma, T., Pedro, K., Cerati, G., Kowalkowski, J., Perdue, G., Spentzouris, P., Tran, N., Vlimant, J.-R., Zlokapa, A., Pata, J., Spiropulu, M., An, S., Aurisano, A., . . . Usher, T. (2020). *Graph neural networks for particle reconstruction in high energy physics detectors* (arXiv:2003.11603). arXiv. https://doi.org/10.48550/arXiv.2003.11603

Krause, C., Winterhalder, R., Feickert, M., & Nachman, B. (2026). *The living guide of machine learning for particle physics* (arXiv:2608.09531). arXiv. https://doi.org/10.48550/arXiv.2608.09531

Leeman, J., Liu, Y., Stiles, J., Lee, S. B., Bhatt, P., Schoop, L. M., & Palgrave, R. G. (2024). Challenges in high-throughput inorganic materials prediction and autonomous synthesis. *PRX Energy, 3*(1), Article 011002. https://doi.org/10.1103/PRXEnergy.3.011002

LIGO Scientific Collaboration, Virgo Collaboration, & KAGRA Collaboration. (2026). Direct multi-model dark-matter search with gravitational-wave interferometers using data from the first part of the fourth LIGO-Virgo-KAGRA observing run. *Physical Review Letters*. https://doi.org/10.1103/qdnd-rsnh

Lookman, T., Liu, Y., & Gao, Z. (2026). Materials informatics: Emergence to autonomous discovery in the age of AI. *Advanced Materials, 38*(29), Article e15941. https://doi.org/10.1002/adma.202515941

Mehta, P., & Schwab, D. J. (2014). *An exact mapping between the variational renormalization group and deep learning* (arXiv:1410.3831). arXiv. https://doi.org/10.48550/arXiv.1410.3831

Merchant, A., Batzner, S., Schoenholz, S. S., Aykol, M., Cheon, G., & Cubuk, E. D. (2023). Scaling deep learning for materials discovery. *Nature, 624*(7990), 80–85. https://doi.org/10.1038/s41586-023-06735-9

Merchant, A., & Cubuk, E. D. (2023, November 29). *Millions of new materials discovered with deep learning*. Google DeepMind. https://deepmind.google/discover/blog/millions-of-new-materials-discovered-with-deep-learning/

Nousi, P., Koloniari, A. E., Passalis, N., Iosif, P., Stergioulas, N., & Tefas, A. (2023). Deep residual networks for gravitational wave detection. *Physical Review D, 108*(2), Article 024022. https://doi.org/10.1103/PhysRevD.108.024022

Okpechi-Kanayochukwu, U. P., Nwaokorogwu, E. C., Ugochukwu, J., & Uzo-Njoku, U. A. (2026). Quantum Biophysical Framework for Modeling Neural Dynamics: Entanglement, Decoherence, Synchronization, and Their Role in Emergent Cognition and Brain Function. *Nigerian Journal of Physics, 35*(S), 442-453. https://doi.org/10.62292/njp.v35(s).2026.766

Panichi, G., Corli, S., & Prati, E. (2026). Quantum physics-informed neural networks for multivariable partial differential equations. *Physical Review Applied, 25*, Article 014001. https://doi.org/10.1103/6nh4-yh2y

Pfau, D., Spencer, J. S., Matthews, A. G. D. G., & Foulkes, W. M. C. (2020). Ab initio solution of the many-electron Schrödinger equation with deep neural networks. *Physical Review Research, 2*(3), Article 033429. https://doi.org/10.1103/PhysRevResearch.2.033429

Physics World. (2024, October 8). *John Hopfield and Geoffrey Hinton share the 2024 Nobel Prize for Physics*. https://physicsworld.com/a/john-hopfield-and-geoffrey-hinton-share-the-2024-nobel-prize-for-physics/

Raissi, M., Perdikaris, P., & Karniadakis, G. E. (2019). Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics, 378*, 686–707. https://doi.org/10.1016/j.jcp.2018.10.045

Royal Swedish Academy of Sciences. (2024, October 8). *The Nobel Prize in Physics 2024* [Press release]. https://www.nobelprize.org/prizes/physics/2024/press-release/

Schlembach, M. C., & Wrublewski, D. T. (2025). Analysis for the science librarians of the 2024 Nobel Prize in Chemistry: Computational protein design and protein structure prediction. *Science & Technology Libraries*, *44*(3), 211–226. https://doi.org/10.1080/0194262X.2025.2468333

ScienceDaily. (2025, January 30). *New diagnostic tool will help LIGO hunt gravitational waves*. https://www.sciencedaily.com/releases/2025/01/250130161650.htm

Sedykh, A., Podapaka, M., Sagingalieva, A., Pinto, K., Pflitsch, M., & Melnikov, A. (2024). Hybrid quantum physics-informed neural networks for simulating computational fluid dynamics in complex shapes. *Machine Learning: Science and Technology*, *5*(2), Article 025045. https://doi.org/10.1088/2632-2153/ad43b2

Szymanski, N. J., Rendy, B., Fei, Y., Kumar, R. E., He, T., Milsted, D., McDermott, M. J., Gallant, M., Cubuk, E. D., Merchant, A., Kim, H., Jain, A., Bartel, C. J., Persson, K., Zeng, Y., & Ceder, G. (2023). An autonomous laboratory for the accelerated synthesis of novel materials. *Nature*, *624*(7990), 86–91. https://doi.org/10.1038/s41586-023-06734-w

Wood, B. M., Dzamba, M., Fu, X., Gao, M., Shuaibi, M., Barroso-Luque, L., Abdelmaqsoud, K., Gharakhanyan, V., Kitchin, J., Levine, D. S., Michel, K., Sriram, A., Cohen, T. S., Das, A., Sahoo, S., Rizvi, A., Ulissi, Z., & Zitnick, L. (2025). UMA: A family of universal models for atoms. In *Advances in Neural Information Processing Systems* (Vol. 38, pp. https://doi.org/129391–129427) https://proceedings.neurips.cc/paper_files/paper/2025/hash/bbf23e81b0ad7637fe9a731d0b676ca6-Abstract-Conference.html

Yu, W., Abdelaleem, E., Nemenman, I., & Burton, J. C. (2025). Physics-tailored machine learning reveals unexpected physics in dusty plasmas. *Proceedings of the National Academy of Sciences*, *122*(31), Article e2505725122. https://doi.org/10.1073/pnas.2505725122

Zeni, C., Pinsler, R., Zügner, D., Fowler, A., Horton, M., Fu, X., Wang, Z., Shysheya, A., Crabbé, J., Ueda, S., Sordillo, R., Sun, L., Smith, J., Nguyen, B., Schulz, H., Lewis, S., Huang, C.-W., Lu, Z., Zhou, Y., . . . Xie, T. (2025).A generative model for inorganic materials design. *Nature*, *639*(8055), 624–632. https://doi.org/10.1038/s41586-025-08628-5