SCIENTIFIC HORIZONS

A Quarterly Review of Discovery & Innovation

Vol. 12 · No. 3 · Autumn 2024

In This Issue

Published by the Scientific Horizons Press | Established 2012
ISSN 2754-0182 | www.scientifichorizons.quarterly

Scientific Horizons

A Quarterly Review of Discovery & Innovation · Vol. 12, No. 3, Autumn 2024

Contact: Scientific Horizons Press, 42 Cavendish Square, London W1G 0AN, United Kingdom. `editorial.quarterly`

From the Editor

The practice of science has always been shaped by its tools. The telescope, the microscope, the particle accelerator — each opened a new window onto nature and, in doing so, transformed not only what we know but how we think. Today, we stand at the threshold of another such transformation, driven not by a single instrument but by a confluence of computational power, algorithmic ingenuity, and data abundance.

In this issue of *Scientific Horizons*, our feature article by Vasquez and Chen examines one of the most exciting frontiers in this new landscape: the application of machine learning to the discovery of quantum materials. The promise is extraordinary — algorithms that can predict stable crystal structures, identify candidates for superconductivity, and even suggest synthetic pathways, all before a single experiment is performed. But as our authors remind us, the algorithmic tail must not wag the scientific dog. Close collaboration between computational and experimental researchers remains essential.

Equally compelling is Nakamura's perspective on the clinical translation of CRISPR-based therapies. A decade after the first demonstrations of programmable gene editing in human cells, we are now seeing treatments move from bench to bedside. The challenges — delivery, specificity, and equitable access — are daunting, but the trajectory is unmistakable.

Finally, our review by Okafor, Sharma, and the CMIP7 Consortium surveys the state of climate modelling at the exascale. As computing power grows, so too does our ability to resolve the fine-scale processes that govern cloud formation, ocean mixing, and ice-sheet dynamics. The picture that emerges is sobering but actionable — a call to both scientific and policy communities.

We hope this issue informs, challenges, and inspires. Science is a collective endeavour, and every reader of these pages is part of the conversation.

— Margaret Okonkwo, Editor-in-Chief

London, September 2024

Contents

DEPARTMENTS

FEATURE ARTICLES

Cover: Crystallographic structure of the perovskite $CaTiO_3$, rendered from neutron diffraction data. Image courtesy of the Materials Project (Lawrence Berkeley National Laboratory).

Design: Elena Petrova. Typeset in TEX Gyre Pagella and Fira Sans.

FEATURE ARTICLE

The Computational Frontier: Machine Learning Meets Quantum Materials Discovery

Elena Vasquez[1] · *James Chen*[2]

In the summer of 2023, a research group at Nexa DeepMind published a landmark study in *Nature* reporting the computational prediction of over 380,000 stable crystalline materials — a number that nearly doubled the cumulative output of a century of experimental crystallography. The tool they used, Graph Networks for Materials Exploration (GNoME), relied on deep graph-neural networks trained on data from the Materials Project, a US Department of Energy initiative that has systematically catalogued the properties of known inorganic compounds since 2011.

This was not an isolated breakthrough. Across condensed-matter physics and materials science, machine-learning (ML) techniques are rapidly reshaping the discovery pipeline. From predicting superconducting transition temperatures to identifying optimal dopants for battery cathodes, data-driven approaches are increasingly complementing — and in some cases supplanting — traditional methods rooted in density functional theory (DFT) and empirical intuition.

Yet the convergence of ML and materials science raises fundamental questions that extend beyond raw predictive power. How do we validate computationally-predicted materials when the volume of predictions outstrips the capacity of experimental verification by orders of magnitude? What does it mean for a crystal structure to be "stable" when the definition depends on the choice of reference chemical potentials? And critically, can ML models ever generate genuine physical insight, or are they merely sophisticated interpolators operating within the boundaries of their training data?

The present article surveys the state of this rapidly evolving field and identifies three grand challenges that will determine whether the marriage of ML and materials science fulfills its considerable promise.

From DFT to Deep Learning

The computational study of materials has, for four decades, been dominated by DFT and the family of methods built upon it. DFT-based high-throughput screening — in which thousands of candidate structures are evaluated against thermodynamic stability criteria — has become a standard workflow, exemplified by the Materials Project, AFLOW, and the Open Quantum Materials Database (OQMD). These databases now contain over two million entries and have enabled the discovery of numerous functional materials, including new thermoelectrics, transparent conductors, and Li-ion conductors.

However, DFT-based screening faces inherent limitations. First-principles calculations are computationally expensive, scaling as $O(N^3)$ with system size, making exhaustive exploration of compositional and configurational space infeasible. Moreover, standard exchange-correlation functionals (e.g., PBE) introduce systematic errors in band gaps, formation energies, and reaction barriers — quantities central to materials design. Hybrid functionals and many-body methods (GW, RPA) improve accuracy but at prohibitive computational cost.

ML models offer a compelling alternative by learning the mapping from structure to property directly from data. Rather than solving the Kohn–Sham equations for each candidate, a trained model can evaluate stability, band structure, or mechanical properties in milliseconds. Early work focused on kernel-based methods (Gaussian processes, support vector regression) trained on hand-crafted descriptors. The field has since shifted decisively toward deep-learning architectures, which can learn representations directly from atomic coordinates and species.

Representation and Architecture

Three classes of neural-network architecture now dominate the materials informatics landscape. Crystal graph convolutional networks (CGCNs) represent a crystal as a graph in which nodes correspond to atoms and edges encode interatomic distances. Message-passing layers aggregate information from neighbouring atoms, recursively building a representation of the local chemical environment. The MatErials Graph Network (MEGNet) and Crystal Graph Convolutional Neural Network (CGCNN) frameworks, both introduced in 2018–2019, remain among the most widely used.

Equivariant neural networks, which explicitly respect the symmetry group of three-dimensional space (E(3)), represent a more recent advance. By constraining their internal representations to transform correctly under rotations and translations, these models achieve state-of-the-art accuracy with fewer training examples. The NequIP and MACE architectures, built on higher-order tensor products of spherical harmonics, have demonstrated remarkable performance on the prediction of interatomic forces and phonon spectra.

Finally, diffusion models and flow-matching approaches, originally developed for image and video generation, are now being adapted to generate crystal structures *de novo*. These generative models learn to reverse a noise-adding process, producing candidate structures that satisfy specified property constraints. GNoME combined a graph-network stability predictor with an active-learning loop, iteratively proposing candidates, evaluating them with DFT, and retraining on the expanded dataset.

The practical impact of these methods is already evident. A 2024 study by Merchant *et al.* used GNoME to identify 52 previously unknown layered materials with potential as solid electrolytes. Independent DFT validation confirmed the stability of 48 of these candidates, and experimental synthesis has been initiated at three laboratories. Similarly, Chen *et al.* reported a graph-network-driven search for high-entropy alloy catalysts that yielded two compositions outperforming commercial Pt/C in oxygen reduction reaction activity.

> *“The algorithmic tail must not wag the scientific dog. Close collaboration between computational and experimental researchers remains essential for realizing the promise of ML-driven discovery.”*

Three Grand Challenges

Despite these successes, the field confronts obstacles that will require concerted effort across the computational and experimental communities to overcome.

Challenge 1: Validation at scale. When a single ML study can predict hundreds of thousands of stable compounds, the bottleneck shifts decisively to experiment. Automated synthesis platforms — including robotic solid-state synthesis and high-throughput thin-film deposition — are beginning to address this gap, but throughput remains orders of magnitude below what computational prediction can produce. The community must develop tiered validation protocols that prioritise candidates with the highest probability of successful synthesis and the greatest potential impact.

Challenge 2: Uncertainty quantification. Most ML models for materials prediction provide point estimates without meaningful confidence intervals. When a model predicts a formation energy of $-1.2\,\mathrm{eV/atom}$, how certain are we that the compound is stable? Bayesian neural networks and ensemble methods offer partial solutions, but rigorous uncertainty quantification (UQ) tailored to materials-specific error modes — configurational entropy, metastable polymorphs, anharmonic effects — remains underdeveloped.

Challenge 3: From prediction to understanding.

The most profound critique of ML in the physical sciences is that it trades understanding for prediction. A model that accurately forecasts superconducting T_c without revealing the underlying pairing mechanism may be useful for engineering but contributes little to fundamental physics. Emerging interpretability techniques — including layer-wise relevance propagation, concept activation vectors, and symbolic regression on learned representations — aim to extract human-comprehensible insights from trained networks, but the gap between these tools and the depth of understanding afforded by traditional theory remains wide.

BY THE NUMBERS

380K	Stable materials predicted by GNoME (2023)
200K	Known inorganic compounds in ICSD (2024)
2.2M	Total entries across Materials Project, AFLOW, OQMD
10^4	Speedup of ML inference vs. hybrid DFT
48/52	GNoME-predicted Li conductors validated by DFT

The Road Ahead

The convergence of ML and materials science is not merely a technological trend but a fundamental shift in how we explore chemical and structural space. The field is evolving from a paradigm of *predict-then-validate* toward one of *co-design*, in which computational prediction, automated synthesis, and characterisation are integrated into closed-loop workflows. Early demonstrations of autonomous laboratories — such as the A-Lab at Lawrence Berkeley National Laboratory, which combines ML-driven experiment planning with robotic synthesis and X-ray diffraction characterisation — point toward a future in which the discovery cycle is compressed from years to days.

Realising this vision will require not only algorithmic advances but also cultural change. Materials scientists and condensed-matter physicists must become fluent in the language of machine learning, while computer scientists must engage seriously with the physics that makes materials prediction a uniquely challenging domain. The journals, funding agencies, and academic departments that structure scientific careers must adapt to reward interdisciplinary work that bridges these historically separate communities.

The stakes are high. From better batteries and more efficient solar cells to room-temperature superconductors and quantum information platforms, the materials we need to address the twenty-first century's most pressing challenges are waiting to be discovered. Machine learning, deployed thoughtfully and in partnership with experiment and theory, may be the tool that helps us find them.

Acknowledgements. E.V. acknowledges support from the US Department of Energy, Office of Basic Energy Sciences (DE-SC0022417). J.C. acknowledges support from the European Research Council (Grant No. 101054458) and the Leverhulme Trust. The authors thank A. Merchant, R. Gomez-Bombarelli, and K. Persson for illuminating discussions.

Competing interests. The authors declare no competing interests.

Data availability. All datasets referenced in this article are publicly available through the Materials Project (materialsproject.org), ICSD, and associated repositories.

PERSPECTIVE

CRISPR 2.0: Precision Gene Editing Enters the Clinic

Sarah Nakamura[3]

When Victoria Gray became the first person to receive an investigational CRISPR-based therapy for sickle-cell disease in 2019, the procedure was rightly hailed as a landmark. Four years later, in December 2023, the US Food and Drug Administration approved Casgevy (exagamglogene autotemcel), the first CRISPR/Cas9-based therapeutic, for both sickle-cell disease and transfusion-dependent beta-thalassemia. The approval, which followed similar authorisation by the UK Medicines and Healthcare products Regulatory Agency, marked the culmination of a journey that began with the discovery of CRISPR adaptive immunity in bacteria and the demonstration, by Doudna and Charpentier in 2012, that the Cas9 nuclease could be programmed for sequence-specific DNA cleavage.

Yet the approval of Casgevy, momentous as it is, represents only the first chapter in the clinical translation of genome editing. A second generation of technologies — base editing, prime editing, and epigenetic editing — is already entering human trials, promising greater precision, broader applicability, and fewer off-target effects than first-generation CRISPR-Cas9.

Beyond Double-Strand Breaks

First-generation CRISPR-Cas9 therapy relies on the introduction of a double-strand break (DSB) at a targeted genomic locus, followed by repair through either non-homologous end joining (NHEJ) or homology-directed repair (HDR). While effective — Casgevy works by disrupting the BCL11A erythroid enhancer to reactivate fetal haemoglobin production — DSB-based editing carries inherent risks, including large deletions, complex rearrangements, and activation of the p53 DNA-damage response.

Base editing, pioneered by David Liu's laboratory at the Broad Institute, sidesteps DSBs entirely. Cytosine base editors (CBEs) and adenine base editors (ABEs) fuse a catalytically-impaired Cas9 nickase to a deaminase enzyme, enabling the direct conversion of C·G to T·A or A·T to G·C base pairs without cleaving the DNA backbone. The first clinical trial of a base editor, BEAM-101 for sickle-cell disease, dosed its first patient in 2023. Early results, presented at the 2024 American Society of Hematology meeting, showed robust HbF induction and no evidence of off-target mutations by whole-genome sequencing.

Prime editing, also developed in Liu's laboratory, extends the editing repertoire further. By fusing a Cas9 nickase to an engineered reverse transcriptase and providing a prime editing guide RNA (pegRNA) that both specifies the target and encodes the desired edit, prime editors can perform all twelve possible base-to-base conversions, as well as small insertions (up to ~44 bp) and deletions (up to ~80 bp), without requiring DSBs or donor templates. The first IND application for a prime-editing therapy — targeting chronic granulomatous disease — is anticipated in 2025.

> *"The approval of Casgevy represents only the first chapter in the clinical translation of genome editing. A second generation of technologies is already entering human trials."*

Delivery Remains the Bottleneck

For all the sophistication of editing modalities, delivery remains the central challenge in translating genome-editing therapies to the clinic. The *ex vivo* approach used by Casgevy — harvesting haematopoietic stem cells, editing them in culture, and reinfusing after myeloablative conditioning — is effective but expensive (Casgevy is priced at $2.2 million per patient in the US) and logistically demanding, requiring specialised manufacturing facilities and extended hospitalisation.

In vivo delivery, in which the editing machinery is administered directly to the patient, would dramatically expand the addressable patient population. Lipid nanoparticles (LNPs), the delivery vehicles that enabled the success of mRNA COVID-19 vaccines, have shown promise for hepatic delivery of CRISPR components. Intellia Therapeutics' NTLA-2001, an LNP-formulated Cas9 mRNA and guide RNA targeting the *TTR* gene for transthyretin amyloidosis, demonstrated durable (>90%) TTR knockdown in Phase I/II trials, representing the first clinical proof-of-concept for systemic *in vivo* CRISPR editing.

For tissues beyond the liver, virus-like particles (VLPs) and engineered adeno-associated viruses (AAVs) are under active investigation, though immunogenicity and payload-size constraints remain significant hurdles.

The Ethical Horizon

As genome-editing technologies become more precise and more accessible, the ethical questions they raise grow sharper. Somatic editing for severe genetic diseases commands broad consensus; germline editing, which would introduce heritable changes, does not. The 2018 announcement by He Jiankui of the birth of genome-edited twins, universally condemned by the scientific community, served as a stark reminder that technological capability can outpace ethical consensus.

The World Health Organization's 2021 governance framework for human genome editing, the International Commission on the Clinical Use of Human Germline Genome Editing's 2020 report, and the 2023 Third International Summit on Human Genome Editing in London have each contributed to a gradually coalescing global framework. Whether that framework can accommodate the accelerating pace of technological development — and whether it can ensure equitable access to therapies that could otherwise deepen health disparities — remains to be seen.

Acknowledgements. S.N. is supported by NIH grant R01-HG012034 and the Harvard Stem Cell Institute. She serves on the scientific advisory board of Verve Therapeutics (uncompensated).

Competing interests. S.N. holds equity in Beam Therapeutics and is a named inventor on patents related to base-editing technologies licensed to Beam Therapeutics through the Broad Institute.

Research Highlights

Astrophysics

JWST Reveals Carbonaceous Dust in the Early Universe. Observations from the James Webb Space Telescope have detected the spectral signature of polycyclic aromatic hydrocarbons (PAHs) in a galaxy at $z = 8.3$, corresponding to an epoch just 600 million years after the Big Bang. The finding, reported by Witstok *et al.* in *Nature Astronomy* (2024), challenges existing models of dust production, which require intermediate-mass AGB stars that would not have had time to evolve by $z > 8$. The authors propose that carbonaceous dust may form in the ejecta of core-collapse supernovae on shorter timescales than previously recognised, a hypothesis consistent with recent nucleosynthetic models. — *Nature Astronomy* **8**, 1124–1131 (2024)

Neuroscience

A Molecular Atlas of the Human Blood–Brain Barrier. A multi-institutional consortium led by the Allen Institute for Brain Science has published a comprehensive single-cell transcriptomic atlas of the human blood–brain barrier (BBB), profiling over 140,000 cells from 17 brain regions across 8 donors. The atlas, described in *Cell* (2024), reveals previously unrecognised heterogeneity in endothelial transport machinery, identifies region-specific vulnerabilities relevant to neurodegenerative disease, and provides a reference dataset for efforts to engineer the BBB for therapeutic delivery. — *Cell* **187**, 3487–3504 (2024)

Climate Science

Accelerating Ice Loss from the Amundsen Sea Sector. A new analysis combining satellite altimetry, ice-penetrating radar, and oceanographic mooring data indicates that ice loss from West Antarctica's Amundsen Sea sector has accelerated by 35% over the past decade. The study, published in *Science Advances* by Davison *et al.* (2024), attributes the acceleration to increased intrusions of warm Circumpolar Deep Water onto the continental shelf, driven by shifts in the Amundsen Sea Low associated with anthropogenic forcing. The findings raise the estimated contribution of the Amundsen sector to global sea-level rise to $0.6\,\mathrm{mm\,yr^{-1}}$. — *Sci. Adv.* **10**, eadk7890 (2024)

Synthetic Biology

De Novo Design of Allosteric Protein Switches. Researchers at the University of Washington's Institute for Protein Design have used deep-learning-based protein design (RFdiffusion and ProteinMPNN) to create synthetic allosteric switches — proteins that change conformation and activity in response to binding a small-molecule ligand — entirely from first principles. The designed switches, reported in *Science* by Anishchenko *et al.* (2024), achieve switching ratios exceeding 50-fold and demonstrate the feasibility of computational design of conditional protein function. — *Science* **385**, 178–185 (2024)

Contributors

Elena Vasquez

Assistant Professor, Department of Materials Science, Stanford University

Elena Vasquez leads the Computational Materials Design group at Stanford. Her research combines first-principles electronic-structure calculations with machine-learning methods to accelerate the discovery of functional materials for energy applications. She received her PhD from MIT in 2018 and was a Miller Research Fellow at UC Berkeley before joining the Stanford faculty in 2021.

Sarah Nakamura

Assistant Professor, Division of Genetics, Brigham and Women's Hospital, Harvard Medical School

Sarah Nakamura is a physician-scientist whose research spans the development of novel genome-editing tools and their translation to clinical application. Her laboratory pioneered several base-editing strategies for haematological disorders. She trained in David Liu's laboratory at the Broad Institute and completed her clinical fellowship in medical genetics at Boston Children's Hospital.

Marcus Okafor

Senior Research Scientist, Met Office Hadley Centre, Exeter

Marcus Okafor leads the High-Resolution Climate Modelling group at the Met Office. His research focuses on the representation of cloud and convective processes in global climate models, with an emphasis on leveraging exascale computing to resolve kilometre-scale atmospheric dynamics. He serves on the CMIP7 Panel and the WCRP Grand Challenge on Clouds and Circulation.

James Chen

Professor of Theoretical Physics, Cavendish Laboratory, University of Cambridge

James Chen holds the Cavendish Chair in Condensed Matter Theory. His group develops and applies machine-learning interatomic potentials and equivariant neural networks to problems in materials physics, with a particular focus on ferroelectrics, multiferroics, and quantum materials. He is a Fellow of the Royal Society and a recipient of the Dirac Medal.

Priya Sharma

Associate Professor, Department of Earth System Science, University of California, Irvine

Priya Sharma's group develops and applies coupled Earth-system models, integrating atmospheric, oceanic, land-surface, and ice-sheet components. Her current research focuses on ice-sheet–ocean interactions in Antarctica and their implications for sea-level projections under high-emission scenarios. She is a lead author on the IPCC Seventh Assessment Report.

Scientific Horizons is published quarterly by Scientific Horizons Press,
42 Cavendish Square, London W1G 0AN, United Kingdom.
ISSN 2754-0182 (Print) · ISSN 2754-0190 (Online)

Typeset in LaTeX using TeX Gyre Pagella and Fira Sans.
Printed on FSC-certified paper by Clays Ltd, Suffolk.