

Artificial Intelligence for Materials Science: Transforming Research Paradigms

Focus Review

Yue Li, Shurui Wang, Jianping Wang, Liu Qian,* and Jin Zhang*

 Cite This: <https://doi.org/10.1021/acs.chemrev.6c00012>


 Read Online

ACCESS |

 Metrics & More

 Article Recommendations

ABSTRACT: The evolution of materials science is undergoing a profound paradigm shift driven by artificial intelligence (AI), transitioning from traditional intuition-driven trial-and-error to an accelerated, data-centric, and algorithmically guided discipline. This review examines this transformation through the lens of the materials discovery workflow, structured around two parallel and complementary trajectories. First, we discuss task-specific AI for materials science. We detail its role in distinct stages of the materials discovery pipeline, including hypothesis generation, experimental planning and optimization, characterization, and knowledge discovery. Second, we explore generalist AI for materials science, designed to handle universal scientific tasks. We examine how these systems advance knowledge representation, enable agentic workflows that orchestrate autonomous laboratories, and facilitate human-AI collaborative reasoning. Finally, we provide perspectives on the future ecosystem of AI for materials science (AI4Mat), outlining the critical challenges and strategic directions that must be addressed to realize the full potential of this evolving discipline.



CONTENTS

1. Introduction	A		
2. Task-Specific AI in Materials Science	C		
2.1. Hypothesis Generation	C		
2.2. Experimental Planning and Optimization	E		
2.2.1. Algorithmic Approaches for Experimental Optimization	E		
2.2.2. Hardware Integration for Experimental Optimization	H		
2.2.3. AI-Powered Synthesis Planning	I		
2.3. Characterization and Data Analysis	J		
2.3.1. Microscopy	J		
2.3.2. Spectroscopy and Scattering	K		
2.4. Knowledge Discovery	L		
2.4.1. Unveiling Governing Laws and Symmetries	L		
2.4.2. From Data to Principles and Descriptors	M		
2.4.3. Explaining Complex Systems and Validating Understanding	M		
3. Generalist AI in Materials Science	N		
3.1. Knowledge Representation and Engineering	N		
3.2. Agentic Workflow	P		
3.3. Human–AI Collaborative Reasoning	R		
4. Outlook and Perspectives	S		
4.1. Data Ecosystems and Automated Infrastructure	S		
4.2. Balancing Generality and Specialization	T		
4.3. Bridging Simulation, Experiment, and Theories	T		
4.4. Collaborative Ecosystems in the Age of AI4Mat	U		
4.5. From AI4Mat to Mat4AI	U		
Author Information	V		
Corresponding Authors	V		
Authors	V		
Notes	V		
Biographies	V		
Acknowledgments	V		
References	V		

1. INTRODUCTION

The evolution of human civilization is inextricably linked to our mastery of matter. From the Bronze Age to the Silicon Era, every leap in civilization has been anchored in the discovery and application of new materials. Traditional materials research has undergone a profound evolution: from intuition-driven trial-and-error to theoretical models based on thermodynamic

Received: January 6, 2026

Revised: April 21, 2026

Accepted: May 5, 2026




Figure 1. Publication trends of AI for Science (AI4S) and Materials Science (AI4Mat). (A) Annual publication volumes from 2015 to 2026. Hatched bars indicate predicted values. (B) Penetration rate, defined as the percentage of AI-related papers within the total publication output of each discipline. AI-related publications were identified by topic search: TS = (“AI” OR “Artificial Intelligence” OR “Machine Learning” OR “Deep Learning” OR “Neural Networks”); “Science” and “Materials Science” domains were each defined by a curated set of Web of Science Categories. 2025 values were annualized from partial-year data; 2026 values were projected using the compound annual growth rate.

and physical laws, and subsequently to the widespread use of computational simulation. While these advances have continually improved efficiency, we remain constrained by the inherent complexity of material systems. Materials are quintessential complex systems, where performance depends on multiscale coupling ranging from Ångström-level atomic arrangements to micron-scale microstructures, and ultimately to macroscopic morphologies. This intricate nonlinear coupling means that slight microscopic variations can drastically alter macroscopic outcomes, leading to long development cycles, high failure rates, and the persistent challenge of balancing precise laboratory control with industrial consistency. Faced with a vast chemical space, estimated to exceed 10^{60} for small molecules alone,^{1,2} traditional reductionist approaches or single-scale simulations are becoming increasingly inadequate.

Against this backdrop, AI for Science (AI4S) has emerged as a new paradigm of scientific research. This represents not merely an improvement in efficiency but a profound methodological revolution.^{3–6} AI technologies, particularly deep learning, possess the unique capability to handle high-dimensional data and capture complex nonlinear correlations, enabling them to transcend the approximation limits of physical modeling and extract mapping relationships directly from data. The power of this approach is evidenced by recent breakthroughs across the natural sciences, where AI has mastered complex systems previously thought intractable. Landmark examples range from AlphaFold’s resolution of the protein folding problem^{7,8} to the control of high-temperature plasma in Tokamak fusion reactors^{9–11} and AI-based weather forecasting models that surpass traditional numerical methods.^{12–14} Given that materials science shares the same fundamental characteristics—high dimensionality, complex many-body interactions, and data sparsity—these successes have ignited a profound expectation within the community: can AI similarly revolutionize the discovery and design of materials? Indeed, this potential is rapidly unfolding, as the academic landscape is being reshaped by an exponential surge in publications applying AI to materials science (AI4Mat) over the past decade. Specifically, annual publications for AI4S and AI4Mat have surged 36-fold and 57-fold over the past decade,

projected to reach nearly 68,000 and 25,000 by 2026 (Figure 1A). Beyond absolute volume, the penetration depth is equally striking: while AI4Mat accounted for a mere 0.14% of materials science papers in 2015, this figure has climbed to 5.74% in 2025 (Figure 1B). Extrapolating this trend paints a transformative picture, the adoption rate is expected to reach 17.28% by 2030 and a staggering 52.01% by 2035, signaling that the AI wave is sweeping through the discipline. This wave spans diverse material classes—from inorganic crystals and energy materials to polymers, soft matter, and molecular systems—and has been surveyed from complementary, domain-specific perspectives in several recent reviews.^{15–17}

From an application-oriented perspective, AI tools employed in materials science can be broadly categorized into two classes: Task-Specific AI and Generalist AI. The broader AI community is currently focused on the explosive development of Generalist AI, driven by advances in Large Language Models (LLMs)^{18,19} and the pursuit of Artificial General Intelligence (AGI).^{20,21} However, for the scientific community, this does not represent a straightforward progression where new AI methods simply replace established methodologies. Instead, the transformation of materials science is unfolding along two parallel and complementary trajectories. Task-Specific AI, technically classified as Artificial Narrow Intelligence (ANI),^{22–24} acts as the pragmatic “solver.” It remains the dominant force for addressing concrete scientific problems, both now and in the foreseeable future, whether facilitating fast potential predictions with DFT-level accuracy or optimizing multidimensional parameters for specific synthesis pathways, leveraging its depth and precision in vertical domains. Conversely, Generalist AI, which represents the path toward AGI, serves as the future “conductor”. Powered by the rise of Foundation Models,^{25,26} these systems are beginning to exhibit universal capabilities in knowledge representation, logical reasoning, and task planning. Generalist AI addresses the challenge of integrating and automating scientific workflows by connecting fragmented toolchains via intelligent agents, pointing toward the ultimate vision of “Self-driving Labs.”^{27–30}

In this review, we provide a comprehensive examination of this dual evolution. Section 2 focuses on the landscape of task-specific AI in materials science, detailing how specialized

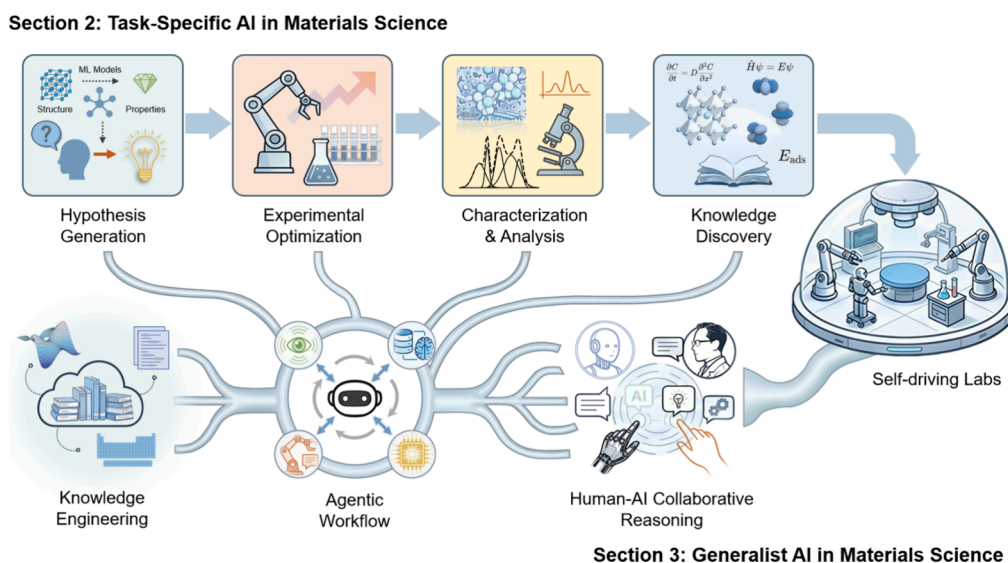


Figure 2. Conceptual framework of parallel AI trajectories in materials science. The top path illustrates Task-Specific AI acting as specialized functional modules (“organs”) along the linear scientific workflow. The bottom path depicts Generalist AI as an integrative network (“blood”) connecting knowledge engineering, autonomous agents, and collaborative reasoning. Just as specialized organs perform vital functions with high specificity, Task-Specific AI modules provide rigorous, domain-grounded predictions at each stage of the workflow; Generalist AI, analogous to the circulatory system, connects these modules by managing the flow of reasoning, context, and instructions across them. This division of labor is analyzed in detail in Section 4.2. Both trajectories converge to enable the ultimate vision of Self-driving Laboratories.

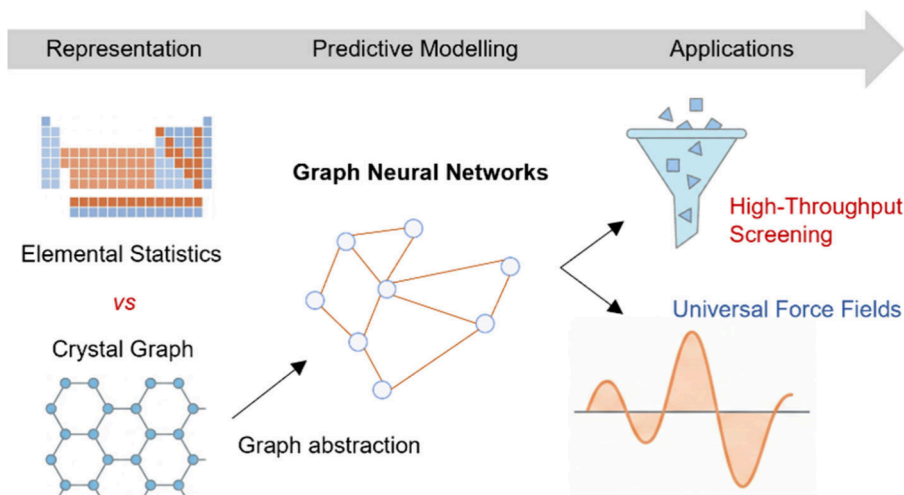


Figure 3. Transition from elemental statistics to crystal graphs enables Graph Neural Networks (GNNs) to serve as the technological base for accurate predictive modeling, driving key applications such as High-Throughput Screening (HTS) and Universal Force Fields.

models are revolutionizing hypothesis generation, experimental optimization, and characterization analysis, and how they are evolving to discover interpretable physical knowledge. Section 3 explores the emerging frontier of generalist AI in materials science, discussing the challenges of knowledge representation and the realization of autonomous agents that can plan and execute scientific tasks, and the evolution of human-AI collaborative reasoning where models serve as cognitive partners (Figure 2). Finally, in this Perspective, we synthesize these developments to provide a perspective on the future ecosystem of AI for Materials Science, discussing the imperative to evolve from data silos to industrialized infrastructure, and how the integration of task-specific AI with generalist AI will bridge the gap between simulation and reality, ultimately advancing the discovery of both new materials and new scientific knowledge.

2. TASK-SPECIFIC AI IN MATERIALS SCIENCE

This section traces how AI addresses specialized scientific problems at each stage of the materials discovery workflow—Hypothesis Generation (2.1), Experimental Planning and Optimization (2.2), Characterization and Data Analysis (2.3), and Knowledge Discovery (2.4)—foregrounding the methodological principles that transfer across material systems.

2.1. Hypothesis Generation

Before the advent of data-driven approaches, the early stage of materials discovery was largely guided by chemical intuition and incremental trial and error. Hypotheses were formulated through heuristic reasoning or extrapolation from limited experimental data, and verification required expensive, time-consuming synthesis or simulations. Only a minute portion of the chemical design space could be explored, and promising but nonobvious candidates often remained undiscovered. In

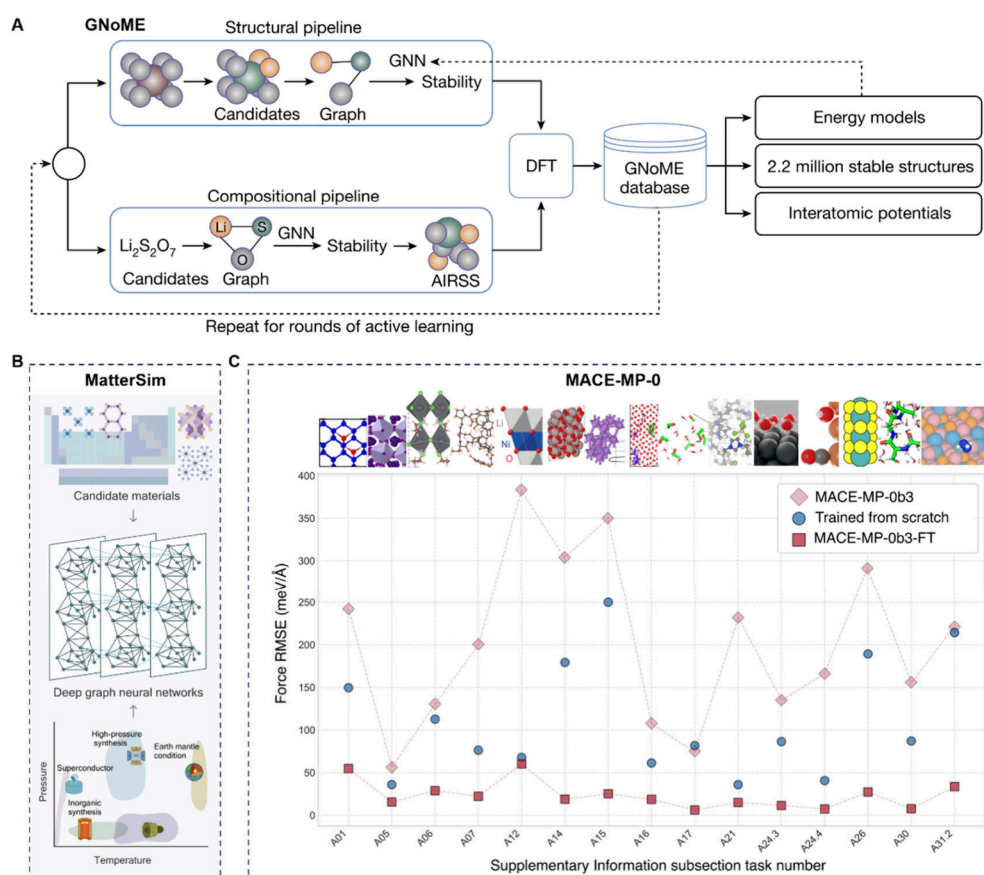


Figure 4. Predictive modeling powering high-throughput virtual screening (HTVS) and machine-learned interatomic potentials. (A) The GNoME workflow as a paradigm for HTVS, utilizing graph neural networks and iterative active learning to significantly expand the known crystal space. Adapted from ref 51. Copyright 2023 under CC BY 4.0 license by the Authors, Springer Nature. (B) MatterSim, illustrating a deep learning potential enabling stable atomistic simulations across extreme thermodynamic conditions. Adapted with permission from ref 55. Copyright 2024 the Authors. (C) Evaluation of MACE-MP-0, a foundation model for atomistic materials chemistry, showing how fine-tuning achieves high accuracy across diverse material systems. Adapted from ref 56. Copyright 2025 under CC BY-NC-ND 4.0 license.

contrast, modern *AI for Science* (AI4S) methodologies replace this empirical bottleneck with algorithmic hypothesis generation and virtual preexperimentation. By learning correlations between structure, composition, and properties, predictive models enable researchers to evaluate millions of hypothetical materials *in silico* before committing to laboratory work.^{22,31,32} This transformation marks the first *forward* step of an AI-driven discovery loop, where ideas are computationally tested prior to synthesis.

Predictive modeling in materials science refers to constructing machine-learning models that learn the “forward” mapping from structure or composition to material properties and serve as computational surrogates for experiments.^{22,32} These models predict thermodynamic stability, electronic structure, mechanical moduli, or catalytic activity directly from digital representations of matter.^{23,33,34} Conceptually, they function as *computational pre-experiments*, acting as rapid estimators that approximate experimental outcomes and guide the allocation of high-value measurements. Within the scientific workflow, predictive modeling provides the quantitative engine for hypothesis generation, high-throughput virtual screening, and the construction of universal interatomic potentials. Its accuracy and transferability depend critically on how materials are represented for learning.

Early approaches relied on manually designed features such as elemental statistics, Coulomb matrices,³⁵ or “sum-over-

bonds” descriptors.³⁶ Although effective for small data sets, these methods required domain expertise and lacked transferability. A paradigm shift occurred with *graph-based representations*, where atoms are nodes and bonds are edges (Figure 3). The Crystal Graph Convolutional Neural Network (CGCNN)²³ demonstrated that graph neural networks (GNNs) can learn property-relevant features directly from atomic connectivity, eliminating the need for handcrafted descriptors. For molecular systems, graph-convolutional architectures achieved superior performance over fixed fingerprints such as Extended-connectivity fingerprints (ECFP).^{37,38} Parallel developments exploited continuous atomic coordinates: the Deep Tensor Neural Network (DTNN)³⁹ and SchNet⁴⁰ introduced continuous-filter convolutions capable of capturing local chemical environments, forming the foundation for modern machine-learning force fields (MLFFs). Composition-only models such as Representation Learning from Stoichiometry (Roost)⁴¹ extended this capability to cases where crystal structures are unknown, enabling property prediction at the earliest stages of design.

Subsequent innovations increased both the physical fidelity and generalization of predictive architectures. MatErials Graph Network (MEGNet)³⁴ incorporated global state variables such as temperature and pressure; improved-CGCNN (iCGCNN)⁴² refined atomic neighborhoods through Voronoi tessellation; and Atomistic Line Graph Neural Network

(ALIGNN)⁴³ explicitly encoded bond angle information via atomistic line graphs. Universal Material Graph with Three-Body Interactions Neural Network (M3GNet)⁴⁴ combined the many-body features of traditional interatomic potentials with flexible graph representations to create a universal potential trained on Materials Project data, achieving broad transferability across 89 elements. To mitigate data scarcity, transfer-learning⁴⁵ and self-supervised pretraining⁴⁶ have become standard strategies, allowing models to leverage knowledge from large DFT corpora or unlabeled structural databases. Adversarial transfer schemes⁴⁷ and physics-embedded architectures⁴⁸ further enhance the extrapolation beyond the training distribution. Collectively, these developments have produced representations that more faithfully reflect the underlying scientific structure of materials, thereby improving the predictive accuracy and enhancing their utility across downstream tasks.

In the discovery workflow, predictive modeling transforms the way hypotheses are formed. Traditionally, scientists begin with a target property—such as high conductivity or catalytic activity—and, guided by physical and chemical reasoning, hypothesize which structural motifs might yield that behavior. In contrast, a predictive model internalizes these structure–property correlations from data, enabling the enumerative generation of hypotheses at scale. Each model prediction effectively constitutes a testable candidate hypothesis within the discovery pipeline—a data-driven assertion that this structure should exhibit such a property, complementing the reasoning- and knowledge-driven hypotheses that arise from domain expertise and physical intuition. When deployed across millions of candidate compositions or structures, the model replaces human intuition with algorithmic exploration, yielding a vast landscape of virtual hypotheses ready for computational or experimental verification. This capability redefines hypothesis generation from a conceptual to a quantitative process, drastically accelerating the early exploratory stage of materials discovery.^{49,50}

These predictive capabilities naturally extend into systematic hypothesis testing through high-throughput virtual screening (HTVS). Fast surrogate models trained on first-principles data can evaluate enormous chemical libraries, ranking candidates according to predicted performance. The GNoME project⁵¹ epitomizes this paradigm: graph neural networks screened over a billion (10^9) candidates, and iterative active learning with DFT verification yielded 2.2 million new stable compounds—expanding the known inorganic crystal space by an order of magnitude (Figure 4A). Similar frameworks have been applied to two-dimensional materials,⁵² perovskites,⁵³ and thermoelectrics.⁵⁴

Beyond screening, predictive modeling underpins the emergence of machine learning potentials that bridge property prediction and atomistic simulation. The M3GNet model⁴⁴ introduced a graph-based, many-body potential capable of near-DFT accuracy across 89 elements. Building on this concept, MatterSim was developed as a deep learning potential actively learned from large-scale *ab initio* data, enabling stable simulations of materials across temperatures from 0 to 5,000 K and pressures up to 1,000 GPa with a 10-fold precision improvement over previous universal MLFFs (Figure 4B).⁵⁵ Most recently, the MACE-MP-0 model established the first foundation model for atomistic materials chemistry, trained on a moderate-sized public data set yet capable of running stable molecular-dynamics trajectories for solids, liquids, gases,

interfaces, and biomolecules alike.⁵⁶ Fine-tuning on minimal additional data allows it to reach *ab initio* accuracy, signaling a major step toward democratizing atomistic simulation (Figure 4C). Although MACE-MP-0 is described as a “foundation model” owing to its broad pretraining across chemical space, its functional role in the discovery workflow remains that of an interatomic potential—predicting energies, forces, and stresses to enable atomistic simulation—placing it within the task-specific trajectory of this section. Collectively, these advances elevate predictive modeling from property estimation to a unified framework for *in-silico* discovery, simulation, and understanding of matter.

Together, these advances establish predictive modeling as the digital front end of the scientific loop. By synergizing the exploratory breadth of HTVS and the physical fidelity of ML potentials, this paradigm transforms hypothesis generation from intuition-driven exploration to an algorithmically guided, quantitative process. The resulting high-confidence insights and candidate rankings produced in this stage form the critical inputs for the subsequent phases of optimization and synthesis planning (Section 2.2), where selected materials are realized experimentally. In essence, predictive modeling transforms the pre-experimental stage from a speculative exercise to a systematic, data-driven search for knowledge and possibility.

Nevertheless, these capabilities come with important caveats. Predictive models are fundamentally constrained by their training data distributions, and their capacity to identify genuinely counterintuitive materials—candidates residing in sparsely sampled regions of chemical space—remains limited. Uncertainty quantification methods, such as ensemble disagreement and Gaussian-process posterior variance, provide a practical safeguard by flagging low-confidence predictions for prioritized verification, but can themselves become unreliable in far-extrapolation regimes. This limitation underscores the necessity of coupling predictive modeling with the iterative experimental feedback strategies discussed in Section 2.2 and the data ecosystem reforms outlined in Section 4.1.

2.2. Experimental Planning and Optimization

The physical synthesis of new materials or molecules remains one of the most resource-intensive bottlenecks in the material discovery workflow. While predictive modeling (Section 2.1) enables the *a priori* estimation of material properties and hypothesis generation, realizing those hypotheses experimentally requires navigating a vast, multidimensional synthesis space governed by uncertain physical and chemical interactions. Traditionally, this process has relied on manual intuition and sequential trial-and-error methods that are slow, expensive, and prone to bias. Artificial intelligence (AI) is now transforming this stage from an intuition-driven search into an algorithmically guided, data-driven optimization process.^{57,58} By coupling learning algorithms with real-time feedback from experiments, AI systems can autonomously plan, execute, and refine synthesis steps, thereby completing the digital-physical loop of discovery that connects virtual prediction to physical realization.

2.2.1. Algorithmic Approaches for Experimental Optimization. Before laboratory automation can reach full autonomy, intelligent algorithms must act as the “computational brain” that decides which experiment to perform next. These algorithms are designed to balance exploration—probing unknown regions of parameter space—and exploitation—refining known promising conditions—while minimiz-

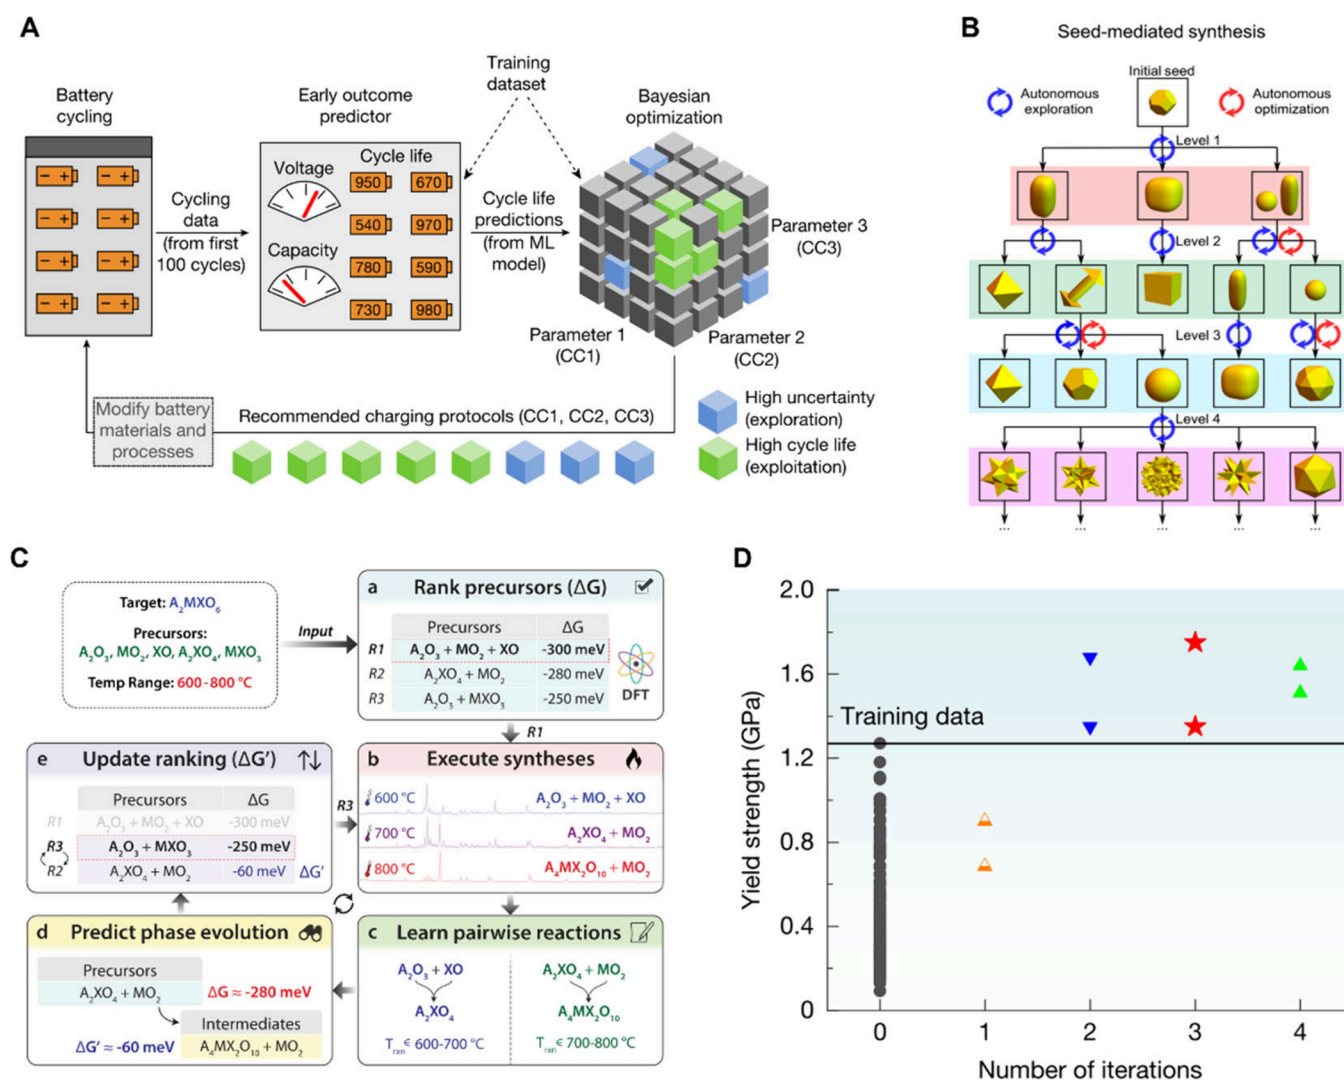


Figure 5. AI-driven experimental optimization strategies across diverse material systems. (A) Closed-loop optimization (CLO) for lithium-ion battery fast-charging protocols. Adapted with permission from ref 62. Copyright 2020 Springer Nature. (B) Heuristic exploration of gold nanoparticle synthesis using the AI-EDISON platform. A Quality-Diversity (QD) algorithm navigates a hierarchical, seed-mediated chemical space to discover distinct morphologies (spheres, rods, stars) and optimize synthesis recipes. Adapted from ref 66. Copyright 2022 under CC BY 4.0 license by the Authors, AAAS. (C) Knowledge-integrated solid-state synthesis planning with ARROWS.³ The workflow combines thermodynamic data with active learning to iteratively rank and select optimal precursor combinations for target materials (e.g., A_2MXO_6). Adapted from ref 70. Copyright 2023 under CC BY 4.0 license by the Authors, Springer Nature. (D) Physics-informed active learning for alloy design. The plot illustrates the rapid discovery of high-strength FeNiCoAlTa alloys (red stars for HEA05), where the algorithm identifies high-performance candidates with significantly fewer iterations compared to random search. Adapted with permission from ref 72. Copyright 2025 Springer Nature.

ing the number of costly experimental trials. Depending on the nature of the design space and available data, approaches can be grouped into probabilistic optimization, heuristic and combinatorial search, generative and reinforcement-learning frameworks, and knowledge-informed strategies that embed domain expertise or physical priors.

Among all optimization strategies, Bayesian Optimization (BO) has proven particularly powerful for continuous experimental variables. A probabilistic surrogate model—often a Gaussian Process—learns the response surface from prior experiments, while an acquisition function selects the most informative next experiment by balancing predicted reward and uncertainty.^{59,60} This sample-efficient strategy is ideal for laboratory conditions where each measurement is expensive. A representative example is the Autonomous Research System (ARES) for carbon nanotube (CNT)

growth.⁶¹ Coupling a BO planner with a chemical-vapor-deposition (CVD) reactor and *in situ* Raman spectroscopy, the system autonomously adjusted temperature and gas flows to maximize growth rate, achieving an 8-fold enhancement over baseline conditions and converging five times faster than a random-forest planner. Similarly, the Closed-Loop Optimization (CLO) framework for lithium-ion batteries employed BO with an early prediction model to design fast-charging protocols (Figure 5A).⁶² By forecasting cycle life from the first 100 cycles, the algorithm identified optimal six-step charging sequences within 16 days instead of >500 days required for exhaustive testing. These systems exemplify how probabilistic reasoning can turn experimentation itself into an iterative learning process. BO has also been applied in atomic-layer deposition,⁶³ polymer formulation,⁶⁴ and organic syn-

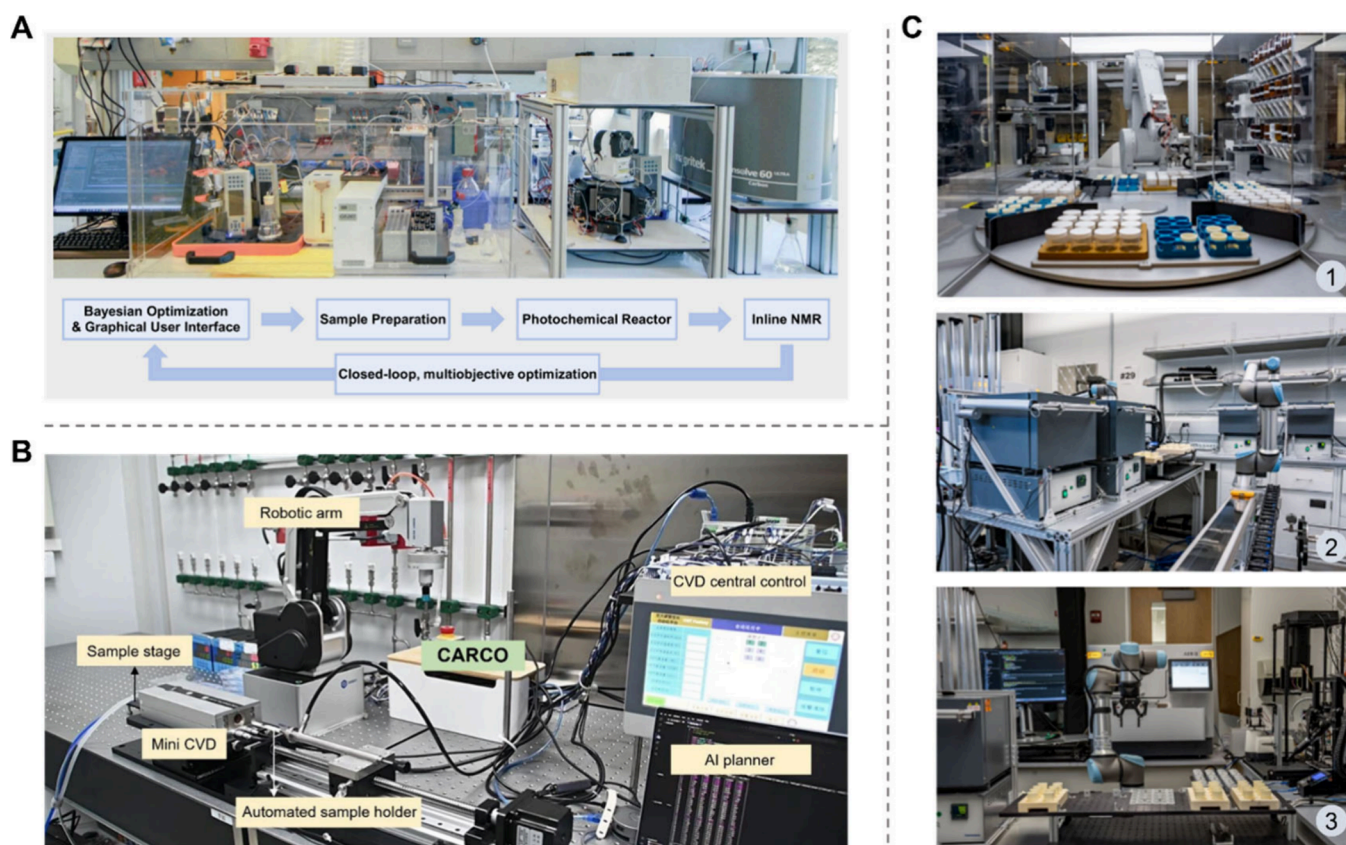


Figure 6. Diverse hardware architectures for automated materials synthesis. (A) Specialized liquid-phase automation. The panel illustrates a continuous-flow platform (RoboChem) integrated with in-line NMR spectroscopy, enabling closed-loop multiobjective optimization of photochemical reactions. Adapted with permission from ref 75. Copyright 2024 AAAS. (B) Specialized multiphase growth hardware. The Carbon Copilot (CARCO) platform automates carbon nanotube synthesis by integrating a robotic arm with a chemical vapor deposition (CVD) system and an AI planner and controller. Adapted from ref 78. Copyright 2025 Elsevier. (C) General-purpose robotics for solid-state synthesis. The A-Lab architecture employs articulated robotic arms to orchestrate discontinuous unit operations, including powder dosing, mixing, and heating, addressing the complex handling requirements of solid reagents. Adapted from ref 85. Copyright 2023 under CC BY 4.0 license by the Authors, Springer Nature.

thesis,⁵⁹ often achieving order-of-magnitude reductions in required experiments.

Many synthesis tasks involve categorical or combinatorial choices—precursor type, solvent, or ligand selection—where gradient-based approaches struggle. Heuristic search algorithms provide a flexible alternative, guided by reward functions or empirical rules that emulate expert reasoning. In peptide-materials design, a Monte Carlo Tree Search (MCTS) framework guided by a random-forest surrogate efficiently explored 3.2 million pentapeptides, discovering unconventional but experimentally validated self-assembling sequences.⁶⁵ In nanomaterials synthesis, the AI-EDISON platform employed a Quality–Diversity (QD) algorithm to navigate the high-dimensional synthesis space of gold nanoparticles (Figure 5B), operating 24 parallel reactors with in-line UV–Vis spectroscopy and discovering five distinct particle morphologies while optimizing nanorod yield to 95%.⁶⁶ Such heuristic and diversity-oriented algorithms emulate the exploratory reasoning of human researchers, mapping entire morphology landscapes rather than searching for a single optimum.

Beyond tuning parameters, the next frontier is inverse design—creating entirely new molecular structures or synthesis recipes that meet predefined objectives. Generative models and reinforcement learning (RL) frameworks achieve this by coupling creative sampling with predictive evaluation in closed

loops. A few-shot inverse-design framework for host-defense-peptide-mimicking polymers combined a multimodal predictive model (text + graph + descriptor) with an RL agent trained through graph-grammar distillation to ensure chemical validity. It screened $>10^5$ virtual polymers and proposed 83 for synthesis, yielding a novel polymer with strong antibacterial activity and low cytotoxicity.⁶⁷ Similarly, the FLAME framework integrated a generative model (REINVENT) with a predictive model (FLSF) trained on 55,000 fluorophore–solvent pairs, facilitating the targeted generation of candidates and leading to the synthesis of a new bright coumarin derivative.⁶⁸ These frameworks extend AI’s role from optimizer to creator, transforming design itself into a learnable, iterative process—one that has proven especially fruitful for structurally modular material classes such as metal–organic frameworks, where the vast combinatorial space of metal nodes, organic linkers, and topologies provides a natural playground for generative approaches.⁶⁹

Integrating scientific knowledge directly into AI-guided optimization can dramatically enhance efficiency and reliability. The ARROWS³ algorithm combined thermodynamic descriptors with active learning to guide precursor selection in solid-state synthesis, identifying all 10 optimal $\text{YBa}_2\text{Cu}_3\text{O}_{6.5}$ routes in 87 iterations—half those required by standard BO—while revealing transferable knowledge about reactive

precursor pairs (Figure 5C).⁷⁰ In the data-driven design of superadhesive hydrogels, data-mining of 24,707 natural adhesive-protein sequences defined an initial search space, which Bayesian optimization then refined to achieve >1 MPa underwater adhesion—an order-of-magnitude improvement.⁷¹ In alloy design, domain-knowledge-informed active learning using physics-based descriptors rapidly identified an FeNi-CoAlTa composition (HEA05) combining 1.8 GPa strength and 25% ductility (Figure 5D).⁷² Optimization accelerates most when algorithms are jointly guided by data correlations and mechanistic understanding. This principle will re-emerge in Section 2.4, where physics-informed AI is examined as a unifying foundation for scientific reasoning.

While these knowledge-informed strategies markedly improve search efficiency, a separate practical concern remains: standard Bayesian optimization relies on assumptions—Gaussian noise, smooth response surfaces, and reproducible measurements—that are frequently violated in real experimental settings. Robust variants employing heteroscedastic noise models and outlier-resistant acquisition functions are beginning to address some of these challenges,⁷³ while human-supervised frameworks that provide active oversight during autonomous campaigns offer a complementary safeguard (discussed further in Section 3.2).⁷⁴

2.2.2. Hardware Integration for Experimental Optimization. Algorithmic intelligence alone cannot revolutionize discovery without physical execution. To complete the discovery loop, predictive and planning algorithms must be seamlessly coupled with robotic hardware, sensing, and automated data acquisition, giving rise to autonomous laboratories—self-driving systems capable of designing, performing, and learning from experiments with minimal human intervention. A closed-loop autonomous laboratory typically integrates four synergistic modules: an algorithmic planner, a robotic executor, automated characterization, and a data-feedback pipeline. In terms of hardware architecture, these systems generally follow two distinct technological paradigms: specialized high-throughput platforms designed for specific tasks, and flexible general-purpose robotics that mimic human manipulation.²⁹

The most established route to hardware automation relies on specialized, task-specific platforms that prioritize throughput and precision for well-defined operations. In the domain of liquid-phase synthesis, this includes both continuous flow reactors and gantry-based liquid handling systems. Continuous-flow platforms, such as the RoboChem system, are inherently automated and allow for the precise control of continuous variables such as residence time and temperature, enabling seamless integration with inline analytics for kinetic studies (Figure 6A).⁷⁵ Similarly, microfluidic platforms have been pivotal in the controlled synthesis of colloidal nanoparticles (e.g., CdSe quantum dots), where automated feedback loops optimize optical properties by adjusting precursor flow rates in real time, achieving throughputs and consistencies unattainable by manual operation.⁷⁶ Complementing flow chemistry, reconfigurable modular platforms have also been developed to handle diverse chemical transformations by switching reactor units, further bridging the gap between specialized synthesis and general applicability.⁷⁷

Beyond liquid-phase synthesis, specialized hardware integration has extended into multiphase growth and device fabrication workflows. In the field of carbon nanomaterials,

the Carbon Copilot (CARCO) platform employs a robotic chemical vapor deposition (CVD) system to automate the synthesis of carbon nanotubes, where growth typically proceeds through vapor–liquid–solid (VLS) multiphase catalytic pathways. By integration of a robotic arm with an automated sample holder and a central controller, CARCO enables the high-throughput screening of catalysts and the precise control of growth parameters (Figure 6B).⁷⁸ Similarly, for two-dimensional (2D) materials, automated systems featuring robotic laminating and delaminating machines have been developed to achieve crack-free, high-fidelity transfer of graphene and MoS₂ wafers on industrial scales.⁷⁹ For optoelectronic devices, platforms such as AMANDA Line One integrate spin-coating, annealing, and spectroscopy modules into a dedicated production line, enabling the optimization of organic solar cells with minimal batch-to-batch variation.⁸⁰ In energy storage, specialized systems such as AutoBASS have been engineered to automate the intricate mechanical assembly of coin-cell batteries, utilizing dedicated grippers and crimpers to ensure high consistency in electrolyte testing.⁸¹

However, complex material research often necessitates a higher degree of flexibility to handle discontinuous unit operations and diverse form factors. This drives the adoption of general-purpose robotics, utilizing articulated robotic arms and mobile robots to connect disparate workstations. Unlike specialized platforms, these systems mimic human researchers and physically transport samples between standard laboratory instruments. A prime example is the Synbot, an AI-driven robotic chemist that autonomously synthesizes organic molecules by orchestrating a fleet of robots across the pantry, dispensing, reaction, and analysis modules, achieving yields comparable to or exceeding human experts.⁸² Similarly, the mobile robotic chemist demonstrated by Burger et al. employs a mobile base to navigate a laboratory, operating benchtop instruments for photocatalyst synthesis and analysis.⁸³ In the domain of solid-state materials, the Ada platform utilizes robotic arms to orchestrate the multistep process of thin-film fabrication, linking synthesis, deposition, and characterization stations to optimize hole-transport materials.⁸⁴ This architecture allows for high reconfigurability, enabling the same robotic hardware to be repurposed for different scientific campaigns.

Despite these advances, a critical bottleneck in hardware integration remains the handling of solid reagents. Unlike liquids, which are easily managed via pumps and pipettes, the automated dispensing and transport of powders require sophisticated weighing and dispensing modules to achieve high precision. Advanced platforms like the A-Lab have begun to address this by integrating powder dosing, ball milling, and high-temperature furnaces into a fully automated workflow for inorganic solid-state synthesis (Figure 6C).⁸⁵ As laboratories continue to scale, ensuring interoperability between these diverse hardware modules—whether specialized gantries or general-purpose arms—becomes paramount. The emergence of open-source hardware standards and communication protocols, such as SiLA 2 (Standardization in Lab Automation), is therefore essential to transition from isolated automated islands to networked, end-to-end discovery ecosystems capable of bridging the gap between digital design and physical reality.⁸⁶ Given the practical difficulty of achieving hardware-level standardization across the diverse instrumentation landscape, software-level abstraction layers offer a complementary and more immediately achievable pathway.

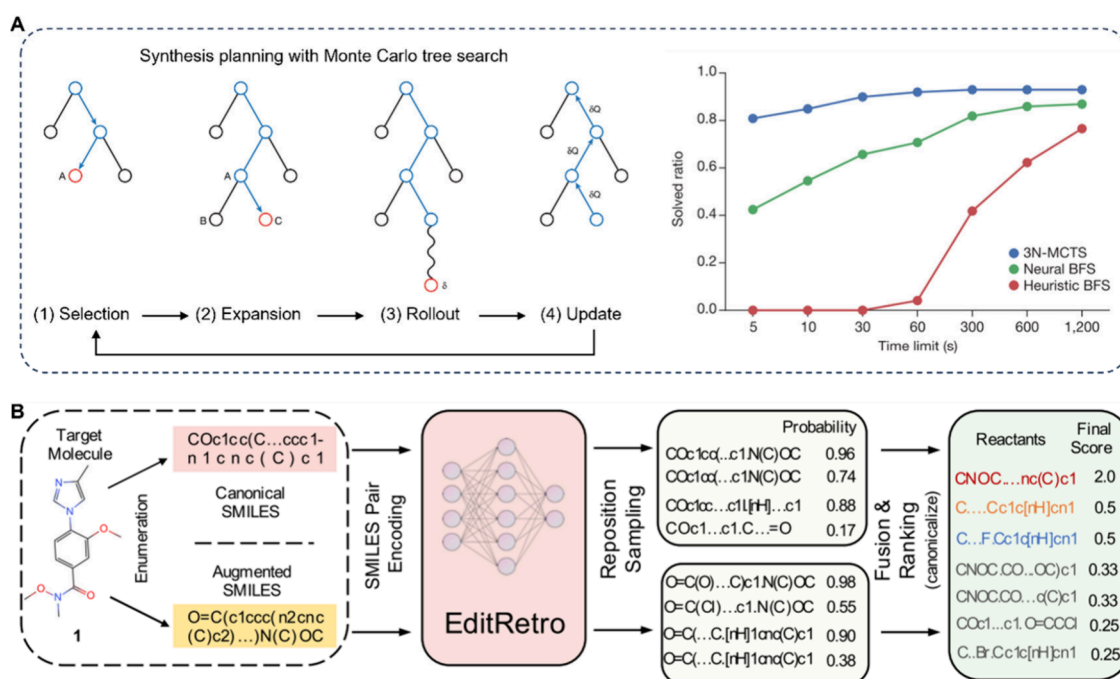


Figure 7. AI-Powered Synthesis Planning. (A) Monte Carlo Tree Search (MCTS) for retrosynthesis. The schematic illustrates the 3N-MCTS framework, which treats synthesis planning as a strategic game involving selection, expansion, rollout, and update steps. The plot on the right demonstrates the superior performance of 3N-MCTS (blue line) in solving complex synthesis routes compared to neural and heuristic breadth-first search (BFS) baselines. Adapted with permission from ref 58. Copyright 2018 Springer Nature. (B) Generative retrosynthesis via structural editing. The workflow depicts the EditRetro model, which represents a conceptual shift by treating retrosynthesis as a sequential structural modification task. The model encodes target molecules (via Canonical and Augmented SMILES) and employs a neural network to predict and rank reactants, mimicking the step-by-step logic of chemical reasoning. Adapted with permission from ref 92. Copyright 2024 Springer Nature.

Analogous to middleware in IoT ecosystems, these layers encapsulate vendor-specific instrument drivers behind unified programmatic interfaces, enabling algorithmic workflows to operate across heterogeneous hardware without modification. A notable example is AEcroscOPy, a cross-platform software-hardware framework that provides a unified Python API for automated and autonomous experimentation across diverse scanning probe and electron microscopy platforms,⁸⁷ effectively decoupling the AI-driven experimental logic from instrument-specific implementation details. Such software abstraction layers may prove more scalable than top-down hardware standardization, particularly for resource-constrained academic laboratories seeking to incrementally automate existing equipment.

2.2.3. AI-Powered Synthesis Planning. Beyond optimizing single steps, synthesis planning determines the entire reaction sequence to construct complex targets. Retrosynthetic reasoning, historically a cognitive task requiring expert intuition, is now automated through machine learning models that learn transformation rules from reaction databases. These systems embody task-level reasoning within chemistry-inferring multistep transformation strategies consistent with known chemical logic.

Foundational work by Coley et al.⁸⁸ utilized neural networks to predict reaction outcomes, providing a quantitative method to evaluate forward feasibility and prevent the generation of chemically impossible routes. Building on this predictive capability, Segler et al.⁵⁸ introduced a landmark hybrid symbolic–neural framework utilizing Monte Carlo Tree Search (3N-MCTS). This approach treated synthesis planning as a strategic game, doubling retrosynthesis success rates over rule-

based baselines and improving planning speed by 30-fold (Figure 7A).

Subsequent progress has moved beyond simple search algorithms to fundamental improvements in how AI perceives and manipulates molecules. To capture subtle reactivity patterns, LocalRetro⁸⁹ integrated local Graph Neural Networks (GNNs) with Transformer attention mechanisms, allowing the model to focus on specific reactive sites within a molecule rather than processing the structure as a generic whole. In parallel, RetroTRAE⁹⁰ adapted the Transformer framework to operate on fragmental and topological descriptors; by decoding local atomic environments instead of text strings, it achieved higher fidelity in chemical mapping. Furthermore, SynCluster⁹¹ introduced reaction-type conditioning to guide pathway generation, effectively clustering strategies based on specific reaction classes to narrow the search space. Representing a conceptual shift, iterative editing models⁹² have recently reframed retrosynthesis as sequential structural modification rather than generation from scratch, effectively mimicking the step-by-step logic of a human chemist (Figure 7B). The integration of these algorithmic advances with human expertise is best exemplified by hybrid systems, such as the computer-aided total synthesis of clovane-sesquiterpenoids.⁹³ In this case, the AI did not merely execute a command but collaborated with researchers to navigate the complex landscape of natural product synthesis. These advances demonstrate that AI can reason within chemical systems—learning implicit reactivity rules and devising plausible synthetic strategies. Yet such reasoning remains domain-specific and goal-bounded. The emergence of agentic AI, models capable of coordinating across prediction, synthesis, and analysis—will extend this

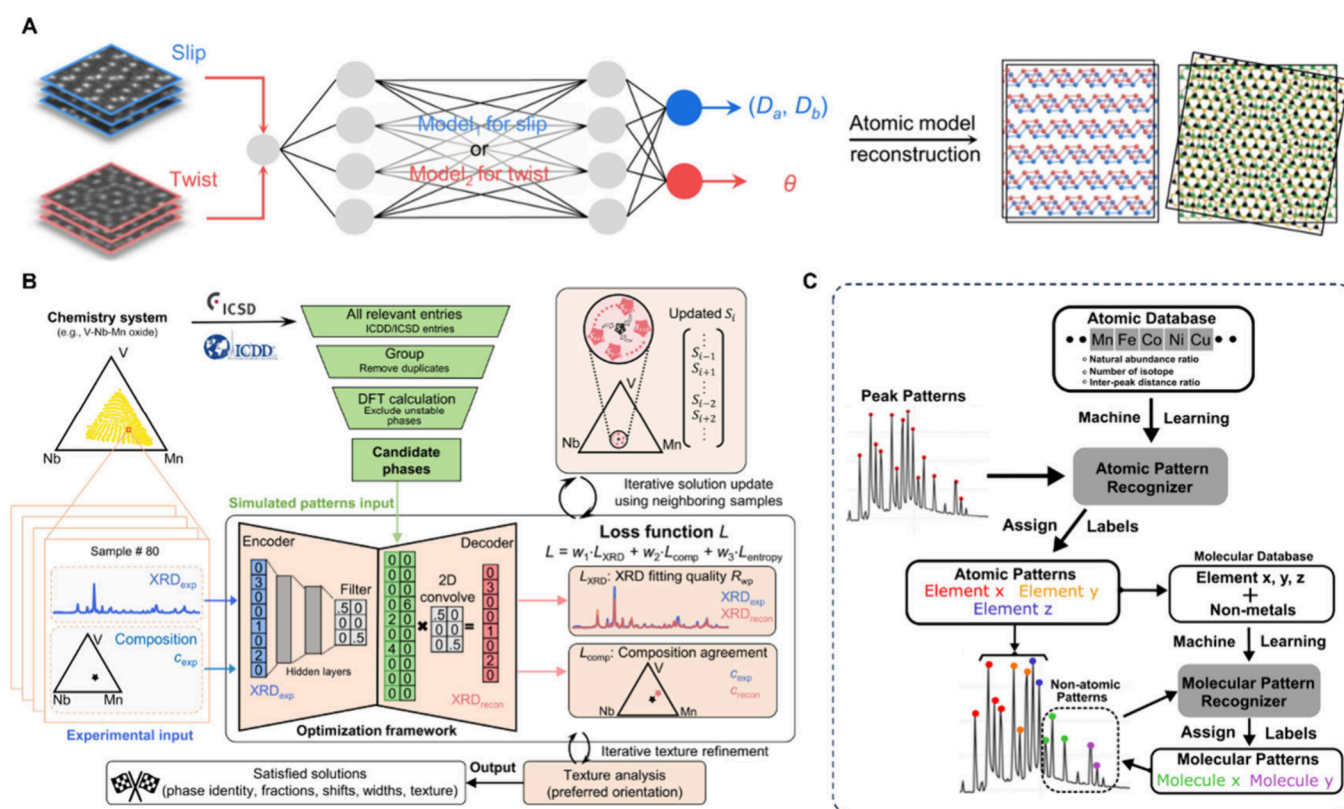


Figure 8. AI-assisted characterization and data analysis. (A) Deep learning inferring interlayer slip and twist angles from STEM images with picometer precision. Adapted with permission from ref 98. Copyright 2025 Springer Nature. (B) Unsupervised phase mapping of high-throughput XRD data encoded with domain-specific knowledge. Adapted from ref 106. Copyright 2025 under CC BY 4.0 license by the Authors, Springer Nature. (C) Automated chemical identification in mass spectrometry via atomic and molecular pattern recognition. Adapted with permission from ref 107. Copyright 2021 Elsevier.

reasoning to the meta-level, a transition explored in Section 3.2.

Optimization and synthesis planning form the operational core of AI-driven discovery. By transforming experiments from static procedures into adaptive learning cycles, AI empowers both the speed and directionality of exploration. Algorithmic intelligence determines how experiments evolve, while robotic platforms execute them; synthesis-planning algorithms connect these steps into coherent strategies. Remaining challenges include data sparsity, interoperability among hardware systems, and standardized feedback protocols. Embedding physics-based priors—briefly touched upon here—offers a route toward robust, interpretable optimization, a concept elaborated in Section 2.4. Together, these developments signify a transition from experiment by intuition to experiment by intelligence, where synthesis itself becomes a self-improving scientific process.

2.3. Characterization and Data Analysis

Optimization and synthesis planning (Section 2.2) depend on how effectively experiments are observed and quantified. In traditional materials research, characterization has long been the rate-limiting and least standardized stage of the discovery cycle. Extracting structural or compositional information from microscopy, diffraction, or spectroscopy data typically requires manual inspection, heuristic fitting, and expert judgment—processes that are time-consuming, subjective, and difficult to reproduce. As instrumentation now generates terabytes of high-dimensional data per day, such manual workflows cannot scale.

AI is transforming this landscape by acting as both an observer and an analyst—automating the conversion of raw experimental signals into standardized, quantitative descriptors of structure, composition, and dynamics. Where Section 2.2 presented AI as the planner that designs and executes experiments, this section positions it as the perceptual interface that grounds those decisions in precise physical observations. Specifically, we examine how AI processes microscopy, scattering and diffraction, and spectroscopic data through feature extraction, denoising, inverse reconstruction, and physics- or multimodal-aware fusion. The analytical toolkit underpinning these pipelines spans a broad methodological continuum—from classical statistical techniques such as principal component analysis (PCA) and non-negative matrix factorization (NMF), which remain indispensable for dimensionality reduction and signal decomposition, to deep learning architectures that extract higher-order representations directly from raw spectral or imaging data. These data-centric capabilities close the loop between synthesis and observation, while Section 2.4 will extend them toward physical interpretation and knowledge discovery.

2.3.1. Microscopy. Modern microscopy techniques (such as TEM, SEM, AFM, and OM) produce vast quantities of complex images. Deep segmentation and detection models now deliver high-throughput, quantitative image analysis with precision and reproducibility beyond manual inspection. For catalytic TEM images, a Cascade Mask-RCNN distinguishes visible nanoparticles from overlapping projections, cutting analysis time from hours to minutes per image while

maintaining an $\approx 2\%$ deviation in mean particle size determination relative to manual analysis.⁹⁴ In diffraction-type imaging, transfer-learned CNNs infer specimen thickness and tilt directly from PACBED patterns in ≈ 0.1 s per pattern—orders of magnitude faster than least-squares fitting with comparable accuracy.⁹⁵

Beyond supervised segmentation, AI is empowering the unsupervised discovery of novel atomic features that might elude human experts. Xue et al.⁹⁶ introduced an unsupervised machine-learning framework combining Zernike polynomials for feature extraction with UMAP-assisted clustering. By automatically processing atomic-resolution STEM images of $\text{MoS}_x\text{Te}_{2-x}$ alloys, the system identified a previously unknown antisite defect configuration ($\text{Te}_{\text{ads-Mo}}$), which was subsequently verified as a highly active catalytic center. This approach demonstrates how AI can extract interpretable structural motifs without the need for laborious manual labeling.

In the realm of imaging resolution, deep learning is bridging the gap between hardware limitations and atomic precision. Qiu et al.⁹⁷ developed SARDiffuse, a denoising diffusion probabilistic model (DDPM) designed to enhance uncorrected STEM images. By learning to restore high-frequency information from experimental aberration-corrected data, SARDiffuse enables subångström spatial resolution (<100 pm) on standard uncorrected microscopes, effectively removing spherical aberration artifacts computationally. This paradigm offers a cost-effective alternative to expensive aberration correctors for high-precision material characterization.

Furthermore, to address the critical bottleneck of data scarcity in training supervised models, generative AI is being deployed to synthesize realistic training data sets. Huang et al.⁹⁸ proposed a “Trident” strategy-enhanced disentangled representation learning method. By generating large volumes of high-quality simulated STEM images that strictly preserve structural labels while mimicking experimental visual styles, they trained a residual neural network to directly infer interlayer slip vectors and twist angles at van der Waals interfaces. This end-to-end approach achieved picometer-level accuracy, resolving complex moiré patterns and stacking registries that are often indistinguishable in the frequency domain (Figure 8A).

Domain-specific transfer learning significantly enhances model generalization in data-scarce regimes. Stuckner et al. demonstrated that encoders pretrained on MicroNet—a repository of over 100,000 labeled microscopy images—outperform generic ImageNet-based models. When applied to nickel-based superalloy segmentation with only a single training image (one-shot learning), the MicroNet-pretrained model reduced the relative intersection over union (IoU) error by 72.2%, highlighting the value of domain-aligned feature representations.⁹⁹ Finally, in situ optical microscopy of MoS_2 CVD growth illustrates bidirectional inference: morphological descriptors extracted from video are linked to synthesis parameters through feature selection and regression, enabling reverse prediction of conditions for large-area crystal growth.¹⁰⁰

Increasingly, these analytical capabilities are being embedded directly into microscopy workflows, enabling self-driving instruments that autonomously direct the data acquisition. AECroscoPy provides foundational infrastructure through a unified Python interface for autonomous experimentation

across probe and electron microscopy platforms.⁸⁷ Built on such frameworks, physics-informed reinforcement learning has achieved autonomous domain-wall design in ferroelectric films,¹⁰¹ while novelty-scoring strategies enable autonomous discovery of unexpected phenomena beyond predefined optimization targets.¹⁰² Complementing these specialized pipelines, the ATOMIC system leverages SAM with ChatGPT-based scheduling for zero-shot 2D material characterization at 99.7% segmentation accuracy, eliminating material-specific training requirements.¹⁰³

2.3.2. Spectroscopy and Scattering. While microscopy visualizes material morphology in real space, spectroscopy and scattering techniques probe internal electronic states and atomic arrangements in energy or reciprocal space. These techniques typically produce high-dimensional signals—ranging from 1D spectra to 3D diffraction volumes—where the key challenge lies in unmixing overlapping signals and identifying hidden fingerprints amidst noise. AI is reshaping this domain by replacing heuristic peak-fitting with robust, high-dimensional feature extraction and probabilistic reasoning.

In the domain of scattering and diffraction, AI automates the recognition of complex structural signatures that traditional heuristics often overlook. Kusne et al. demonstrated this with the CAMEO platform, which integrates Bayesian active learning to autonomously interpret X-ray diffraction (XRD) patterns in real time, enabling the rapid phase mapping of Ge–Sb–Te phase-change materials minimizing human intervention.¹⁰⁴ Expanding on this, Salgado et al. proposed a generalized deep learning framework for the automated classification of crystal systems and space groups from XRD patterns by synthesizing augmented data and optimizing a no-pooling CNN architecture, which reliably deciphers complex structural features often obscured in experimental data, demonstrating strong generalization capabilities on diverse experimental data sets.¹⁰⁵ Moving beyond purely data-driven recognition, AutoMapper, developed by Yu et al., enables rapid and precise analysis of complex XRD patterns, uncovering previously overlooked phases and enabling texture analysis. Its strong performance stems from the encoding of domain-specific knowledge by a designed loss function grounded in crystallographic principles and X-ray diffraction physics (Figure 8B).¹⁰⁶

In parallel, mass spectrometry and spectroscopy have benefited from algorithmic advancements in identifying chemical signatures from noisy backgrounds. Wei et al. introduced a LightGBM-based pipeline for time-of-flight mass spectrometry (ToF-SIMS) that exploits isotopic abundance and interpeak spacing to identify atomic and molecular ions within microseconds, replacing slow, user-dependent annotation with high-throughput precision (Figure 8C).¹⁰⁷ Jia et al. applied non-negative robust PCA—a classical statistical decomposition technique—to electron energy-loss spectroscopy (EELS) imaging, demonstrating how established methods gain renewed utility when integrated into AI-driven analytical workflows. This method decomposes data into low-rank intrinsic signatures and sparse noise, improving the signal-to-noise ratio (SNR) by an order of magnitude and successfully disentangling overlapping Ru/O edges that defy conventional PCA or NMF.¹⁰⁸ Furthermore, ensemble frameworks like ELSIE, developed by Zheng et al., demonstrate the power of simulation-based matching. By comparing experimental inputs against a database of over 800,000 computed

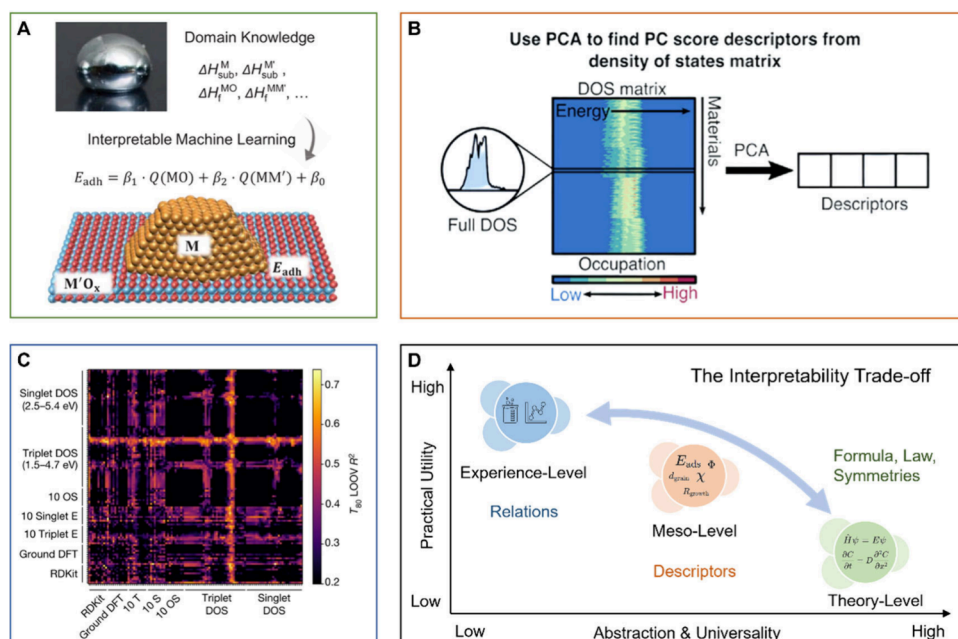


Figure 9. AI-driven knowledge discovery and physical interpretation. (A) Discovery of governing laws via symbolic regression for metal–support interactions. Adapted with permission from ref 112. Copyright 2024 AAAS. (B) Unsupervised extraction of electronic descriptors using PCA on density of states. Adapted with permission from ref 115. Copyright 2021 Elsevier. (C) Explainable modeling identifying triplet density of states as a key predictor for molecular photostability. Adapted with permission from ref 120. Copyright 2024 Springer Nature. (D) The interpretability trade-off, illustrating the hierarchy from empirical relations to universal laws.

spectra, ELSIE accurately identifies oxidation states and coordination environments in X-ray absorption near-edge structure (XANES) spectra, achieving 84.2% accuracy.¹⁰⁹

Collectively, these data-centric capabilities—from automated microscopy analysis to spectral fingerprinting—mark a fundamental shift in materials characterization, though the transferability of most models beyond the specific instruments and laboratories on which they were trained remains an open challenge (Section 4.1). By converting raw experimental signals into standardized, quantitative descriptors, AI closes the loop between synthesis (Section 2.2) and observation. However, extracting parameters is only the first step; the ultimate goal is to uncover the physical laws governing these observations. This transition—from identifying what structure is present to understanding *why* it forms—forms the basis of the next section (Section 2.4), where we explore AI’s role in knowledge discovery and physical interpretation.

2.4. Knowledge Discovery

The preceding sections have illustrated how artificial intelligence accelerates materials research by enhancing prediction, optimization, and data analysis. Yet, these advances—though transformative—remain largely instrumental: AI serves to improve efficiency and performance. This section marks a conceptual shift. Here, the objective is not to achieve higher accuracy or faster discovery, but to pursue understanding—to use AI as a means of extracting interpretable physical knowledge, discovering governing principles, and formulating data-driven laws. In this role, AI becomes a theory builder, transforming from a computational assistant into a participant in scientific reasoning itself.

Unlike the task-specific applications of physics-informed models in earlier sections, which use knowledge to enhance prediction or optimization, the works discussed here employ knowledge integration to generate new knowledge. They seek

to formalize mechanisms, expose symmetries, and construct compact, physically grounded relationships that bridge the gap between empirical observations and fundamental understanding. These efforts can be viewed across three tiers: the law-level formalization of governing eqs (Section 2.4.1), the meso-level discovery of empirical rules and descriptors (Section 2.4.2), and the system-level explanation of complex, high-dimensional behaviors (Section 2.4.3).

2.4.1. Unveiling Governing Laws and Symmetries. AI can elevate materials science from empirical heuristics to law-level formalisms that compactly express how systems behave across chemical and structural families. Among the most direct approaches is symbolic regression, which identifies the minimal mathematical expressions capable of reproducing observed dynamics. The Sparse Identification of Nonlinear Dynamical systems (SINDy) framework exemplifies this idea: by constructing a large library of candidate functions and selecting only the few necessary to reproduce measured time derivatives, it recovers governing equations directly from data.¹¹⁰ Applied to canonical problems such as the Lorenz attractor and fluid vortex shedding, SINDy rediscovered the correct underlying differential equations—compressing decades of manual model building in fluid dynamics into a single computational workflow. Building on this, the AI Feynman algorithm integrates physics-inspired strategies—such as dimensional analysis and symmetry detection—to rediscover fundamental physical laws from noisy data, demonstrating a capability to derive exact formulas from the Feynman Lectures on Physics.¹¹¹ In a materials context, this ability to infer closed-form evolution laws from transient experimental or simulation data opens a route to mechanistic rate equations for phase transformations, defect diffusion, or thin-film growth, providing interpretable and testable expressions rather than opaque predictive fits.

A striking example of such AI-discovered law in materials chemistry is the work of Wang et al. on metal–support interactions (MSIs) in heterogeneous catalysis. Using advanced symbolic regression trained on 178 metal/oxide interfaces, the authors derived an analytic two-term expression for the interfacial adhesion energy. This concise law unifies a wide range of interfacial chemistries and reveals a new physical principle: for late transition metals, it is the strength of metal–metal interactions—rather than traditional oxophilicity—that governs encapsulation behavior in strong metal–support interactions (SMSI). The work demonstrates how AI can move beyond correlation to codify mechanistic principles in analytic, generalizable form (Figure 9A).¹¹²

Complementary to explicit equation discovery, machine learning can reveal hidden symmetries—the latent coordinate transformations under which physical laws take their simplest form. The framework proposed by Liu and Tegmark parametrizes all possible invertible coordinate mappings as a neural network and defines “asymmetry” as deviation from a target partial differential equation that encodes a physical invariance. By minimizing this asymmetry loss, the model identifies the coordinate system in which the symmetry is most apparent. Although demonstrated on canonical physical systems—such as recovering the Gullstrand–Painlevé transformation of the Schwarzschild metric or identifying Hamiltonian structure in coupled oscillators—the implications for materials science are tangible. In high-dimensional representations of microstructures, spectra, or reaction networks, discovering such transformations can linearize complex couplings and expose the low-dimensional manifolds on which materials phenomena truly evolve.¹¹³

Together, these studies mark a transition from AI as a predictive model to AI as a law-discovery mechanism, one that formulates the compact mathematical statements through which new material principles can be articulated.

2.4.2. From Data to Principles and Descriptors. Moving from universal laws toward mesoscopic regularities, AI enables the discovery of interpretable descriptors and phenomenological principles that generalize across related systems without claiming absolute universality. In this intermediate regime, models act as quantitative microscopes, revealing the variables, thresholds, and failure modes that govern structure–property relationships.

A leading paradigm in this domain is the sure independence screening and sparsifying operator (SISSO), which tackles the immense combinatorial space of potential physical parameters. By iteratively combining primary features (e.g., atomic radii, ionization energies) into candidate expressions and selecting the most predictive yet concise ones, SISSO has successfully identified low-dimensional descriptors for predicting the phase stability and classifying the metallicity of binary compounds.¹¹⁴

Complementary to symbolic methods, unsupervised learning can extract electronic descriptors directly from high-dimensional quantum data. Esterhuizen et al. applied PCA to large databases of *d*-projected density of states, and found that the first two principal components correspond naturally to the *d*-band center and width, while subsequent components capture filling and skewness. These automatically learned axes reconstruct established chemical intuition—linking geometry and composition to chemisorption energy—but without relying on predefined features. Such unsupervised learning thereby rederives the electronic “periodic table” of bonding tendencies directly from data (Figure 9B).¹¹⁵

AI can also illuminate where accepted empirical rules fail and why. Using subgroup discovery and outlier analysis on a database of nitrogen-reduction catalysts, Li et al. identified specific subclasses—early transition-metal single atoms on VS₂—where the conventional linear scaling between N and NH₂ adsorption energies breaks down. Subsequent electronic-structure analysis traced this anomaly to strong charge transfer from the metal site to the support. In this way, AI does not merely refine existing relationships but maps the boundaries of their validity, guiding human researchers to the regions of parameter space where new physics emerges.¹¹⁶

In sum, meso-level discoveries—be they latent electronic descriptors, symbolic design rules, or identified rule-breaking regimes—serve as the connective tissue between data and theory. They provide interpretable, transferable, and diagnostic insight, transforming raw numerical patterns into the conceptual building blocks of future physical laws.

2.4.3. Explaining Complex Systems and Validating Understanding. At the system level, explainable and knowledge-integrated AI demonstrates how these principles translate into causal understanding and experimentally verifiable insight. Here, interpretability becomes not a byproduct but a primary objective—bridging high-dimensional models with human reasoning to validate and refine scientific understanding.

When interactions are high-order and nonlinear, explainable AI (XAI) clarifies structure–property and process–performance relationships while preserving predictive strength. In the floating-catalyst CVD growth of carbon nanotubes, Bulmer et al. trained a multivariate model on 631 experiments to quantify the relative importance of CO₂ flow rate, water vapor, and temperature, revealing nontrivial interactions between oxidative species and thermal conditions. These data-driven relationships were then rationalized through a kinetic–thermodynamic model of catalyst-particle evolution, demonstrating a full causal chain from empirical trends to physical mechanism.¹¹⁷

Across diverse materials domains, SHapley Additive exPlanations (SHAP) analysis has proven especially powerful for extracting such quantitative insight. Kudhanya and Lyth used SHAP to resolve longstanding contradictions in hydrogen physisorption on porous carbons, identifying the total pore volume and oxygen content as dominant positive contributors.¹¹⁸

Beyond feature attribution, interpretability extends to uncertainty and corrigibility. Feng et al. constructed a probabilistic graphical model for oxygen-reduction reaction microkinetics that decomposes uncertainty into distinct physical sources—DFT error, solvation error, and data correlations—producing an “uncertainty blueprint” that quantitatively indicates which components of the model require refinement. Such frameworks shift the focus from predicting outputs to understanding confidence, thereby guiding future theory correction.¹¹⁹

Finally, closed-loop transfer systems integrate interpretable modeling directly into autonomous experimentation. In a landmark study on donor–acceptor molecules, Angello et al. combined Bayesian optimization for experiment planning with interpretable support vector regression on time-dependent DFT descriptors. The resulting human–AI hybrid workflow revealed that triplet density of states—rather than the lowest triplet level—governs molecular photostability, a counterintuitive principle experimentally confirmed through targeted

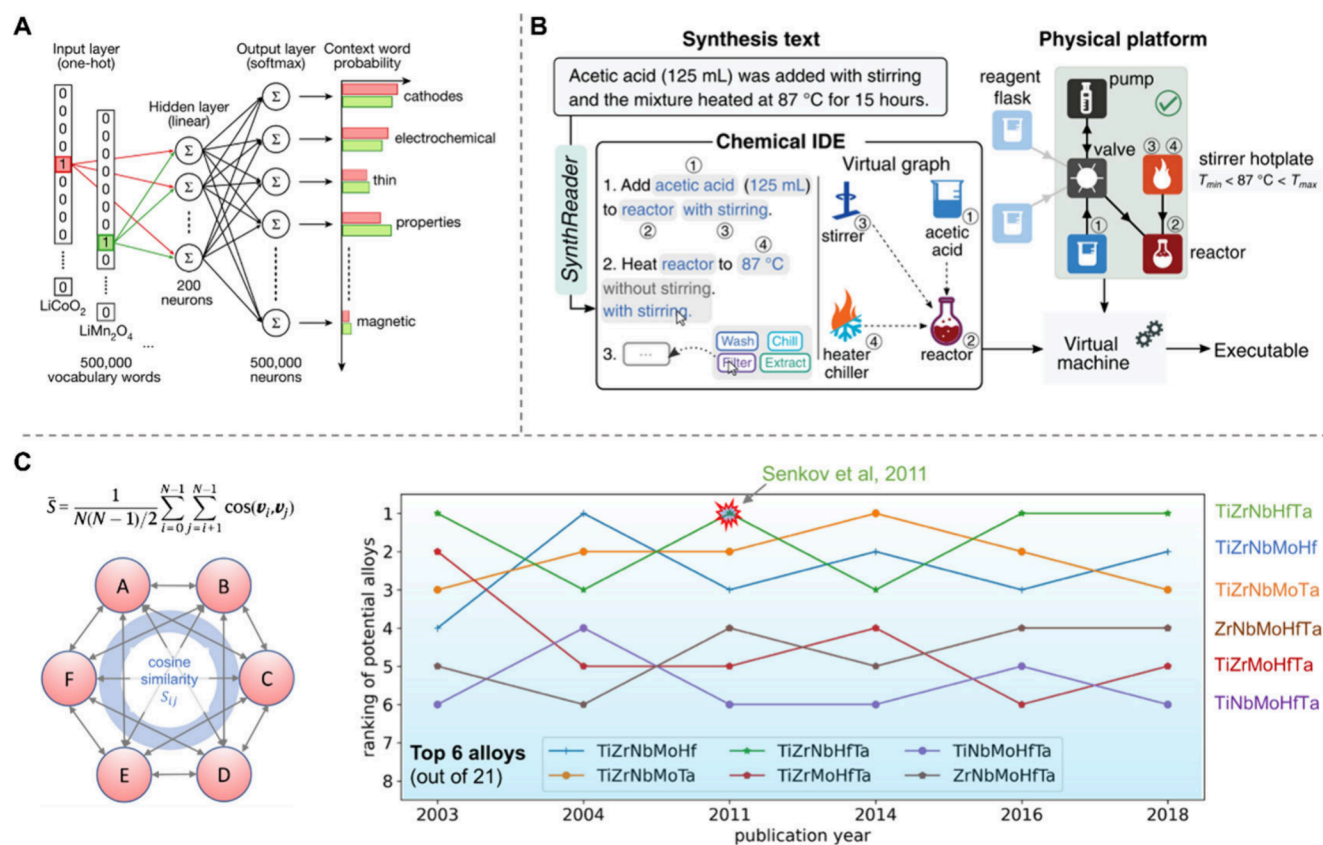


Figure 10. Knowledge representation and engineering in materials science via Generalist AI. (A) Unsupervised latent embeddings. A Word2Vec-based neural network architecture maps scientific concepts into dense vector spaces, capturing implicit chemical analogies and functional properties from unstructured text. Adapted with permission from ref 121. Copyright 2019 Springer Nature. (B) Structured execution via ontologies. The Chemical Description Language (XDL) framework translates synthesis procedures into hardware-agnostic executable graphs, bridging the gap between literature text and robotic automation. Adapted with permission from ref 136. Copyright 2020 AAAS. (C) Physics-aware context similarity. Context metrics derived from large-scale literature mining successfully identify stable high-entropy alloy candidates (e.g., TiZrNbHfTa), linking textual co-occurrence with thermodynamic feasibility. Adapted from ref 138. Copyright 2023 under CC BY 4.0 license by the Authors, Springer Nature.

synthesis (Figure 9C). These examples collectively demonstrate that explanation is not merely a narrative add-on: it is an actionable component of discovery that validates and expands physical understanding.¹²⁰ Nevertheless, the community has yet to establish consensus validation criteria for distinguishing genuine physical insight from statistical artifact in AI-derived laws—a challenge that demands layered assessment spanning statistical robustness, physical consistency, and independent experimental falsification.

The studies summarized above position AI not only as a computational accelerator but as a partner in theory construction—one capable of extracting sparse laws, discovering latent coordinates, distilling mechanistic descriptors, and quantifying uncertainty within physical models. However, the pursuit of physical interpretation involves a fundamental trade-off between practical utility and theoretical universality (Figure 9D). While “Theory-Level” AI distills abstract laws with broad applicability (Section 2.4.1), “Experience-Level” and “Meso-Level” models provide the concrete relations and descriptors necessary for specific material challenges (Sections 2.4.2 and 2.4.3). Crucially, the knowledge extracted in this stage does not merely conclude the workflow—it actively seeds the next iteration of discovery. Descriptors identified through symbolic regression (Section 2.4.1), such as the metal–support interaction parameter derived by Wang et al., can be directly

incorporated as features in subsequent high-throughput virtual screening campaigns (Section 2.1), thereby constraining the search space with physically meaningful criteria. Similarly, SHAP-derived feature attributions (Section 2.4.3) reveal which compositional or structural variables most strongly influence target properties, providing experimentally testable hypotheses that guide the next round of synthesis optimization (Section 2.2). In this way, Section 2 constitutes not a linear sequence but a self-reinforcing cycle in which each stage generates inputs for the others. The next section turns toward a parallel trajectory, Generalist AI, which seeks to integrate these diverse levels of reasoning and tooling into unified, agentic systems capable of navigating the full scientific workflow.

3. GENERALIST AI IN MATERIALS SCIENCE

3.1. Knowledge Representation and Engineering

The transition from task-specific AI to generalist systems in materials science depends fundamentally on how scientific knowledge is represented. Before any model can reason or plan, the vast, multimodal, and often unstructured body of materials data and literature must be transformed into a machine-interpretable form. Over the past decade, this transformation has evolved from unsupervised embeddings that encode latent semantics to structured ontologies that can drive autonomous laboratories. These developments collec-

tively form the cognitive substrate upon which higher-level agentic intelligence can operate.

This substrate has been assembled in four broadly complementary stages—each extending, rather than replacing, its predecessors. The section first surveys unsupervised embeddings (Word2Vec, Atom2Vec, Mol2vec) that capture implicit chemical regularities, then domain-adapted transformers (MatSciBERT, MaterialsBERT, ElementBERT) that add supervised task performance, followed by structured extraction pipelines that convert prose into explicit, queryable records, and finally standardized ontologies and description languages (XDL) that make knowledge executable. The section closes by examining how these representations jointly enable cross-domain reasoning and by identifying multimodality and domain anchoring as persistent bottlenecks.

The earliest attempts to represent materials knowledge relied on unsupervised embeddings trained directly on scientific text and chemical formulas. A landmark demonstration by Tshitoyan et al. showed that a Word2vec model trained on 3.3 million materials abstracts could recover periodic-table organization, chemically valid analogies, and even anticipate promising thermoelectrics years before their experimental discovery (Figure 10A).¹²¹ Analogous ideas were extended from words to atoms and molecular fragments: Atom2Vec learned element embeddings that cluster according to chemical groups and improve downstream property prediction,¹²² while Mol2vec captured chemically meaningful relations among molecular substructures to predict compound properties and bioactivities.¹²³ Building upon this foundation, a blank-filling language model for crystal formulas (BLMM) learned a “chemical grammar” that respects charge neutrality and substitution rules, enabling generative and tinkering design directly in composition space.¹²⁴ These works collectively illustrate that latent vector spaces can internalize chemical regularities and encode domain knowledge without explicit supervision.

Building on these foundations, domain-adapted transformer models extended linguistic representation into a full-fledged scientific language for materials. MatSciBERT, initialized from SciBERT and retrained on a curated corpus of materials publications, outperformed SciBERT on document classification, named-entity recognition, and relation extraction tasks by capturing the field’s specialized semantics.¹²⁵ MaterialsBERT, trained on 2.4 million materials abstracts, expanded these capabilities to polymers and other functional systems, obtaining ~ 300,000 material property records.¹²⁶ Moving toward domain-specific semantic modeling, ElementBERT, a BERT-based language model pretrained on 1.29 million alloy-related abstracts, learned element-level semantic embeddings that encode latent chemical knowledge, enabling improved prediction of thermodynamic and mechanical properties across shape memory alloys (SMAs), Ti alloys, and high-entropy alloys (HEAs), enhancing prediction accuracy by 23% relative to traditional methods, and also consistently outperforming general BERT variants.¹²⁷ Such models begin to blur the line between language and physics, showing that contextual transformer representations can learn oxidation-state constraints, periodic trends, and substitution rules purely from textual co-occurrence.

While latent embeddings encode implicit relations, a parallel effort has focused on converting unstructured scientific prose into explicit, structured records. Early NLP pipelines combined Bi-LSTM-CRF architectures with rule-based templates to

recognize materials entities and their roles in synthesis paragraphs, enabling meta-analyses such as precursor similarity and substitution in solid-state reactions.¹²⁸ Self-supervised “masked precursor completion” further learned synthesis-specific embeddings that outperformed general text or composition representations for recommending precursor sets for unseen targets.¹²⁹ With the advent of large language models, structured extraction has become both simpler and more accurate. LLMs trained on predefined JSON schemas extracted structured data for diverse subfields with F₁ scores reaching 0.82, surpassing BERT-based baselines (MatBERT and seq2rel) on nested and relational tasks,¹³⁰ while fine-tuned sequence-to-sequence LLMs reconstructed complete multi-stage protocols (e.g., seed-mediated gold nanorods) with high fidelity.¹³¹ Rule-augmented systems remain important for scale: automated pipelines have extracted composition–property pairs for superalloys from tens of thousands of articles, including nontrivial table linking, and trained predictive models whose outputs were experimentally verified;¹³² similar text-mined databases exist for magnetic and superconducting transition temperatures¹³³ and MOF synthesis routes (DigiMOF).¹³⁴ Document-embedding approaches (e.g., Doc2vec with UMAP and HDBSCAN) extend these concepts to the level of topic mapping, revealing quantitative evolution of research trends and interdisciplinarity across the materials community.¹³⁵

Structured representation reaches its fullest potential when codified as standardized ontologies and description languages. The Chemical Description Language (χ DL/XDL) expresses syntheses as hardware-agnostic operations (“Add,” “Heat,” “Filter”), allowing literature procedures to be compiled into robotic instructions (Figure 10B).¹³⁶ Recent integration of LLMs as parsers and validators for χ DL has enabled fully autonomous execution on Chemputer platforms, closing the loop from natural-language text to chemical action.¹³⁷ These developments transform representation from descriptive to prescriptive—knowledge is no longer merely stored but executed, advancing reproducibility and automation in scientific practice.

Once representations are learned and structured, they enable reasoning and generalization across domains. Mining 6.4 million abstracts to construct element context vectors yields a “context similarity” metric that recovers canonical high-entropy-alloy systems and correlates with thermodynamic rules for solid-solution formation, bridging text-derived and physics-based knowledge (Figure 10C).¹³⁸ Latent spaces trained on literature have prioritized thermoelectric materials years ahead of experiment,¹²¹ while foundation-model embeddings power multitask recommenders for materials retrieval and design. Large language models tested on the MaScQA benchmark demonstrate competence comparable to average undergraduates but still fail on conceptual nuances, highlighting the need for domain anchoring.¹³⁹ In chemistry, fine-tuned or prompted LLMs predict reaction outcomes and linker mappings directly from text representations (e.g., SMILES, IUPAC), showing that linguistic knowledge can encode functional chemistry.¹⁴⁰ A persistent bottleneck is multimodality: key data are often distributed across text, tables, and figures, requiring future systems that jointly parse and align heterogeneous sources to avoid fragmentation.¹⁴¹

Across these developments emerge four principles: (1) self-supervision on large corpora yields latent spaces that capture chemical regularities; (2) domain-adaptive pretraining enhan-

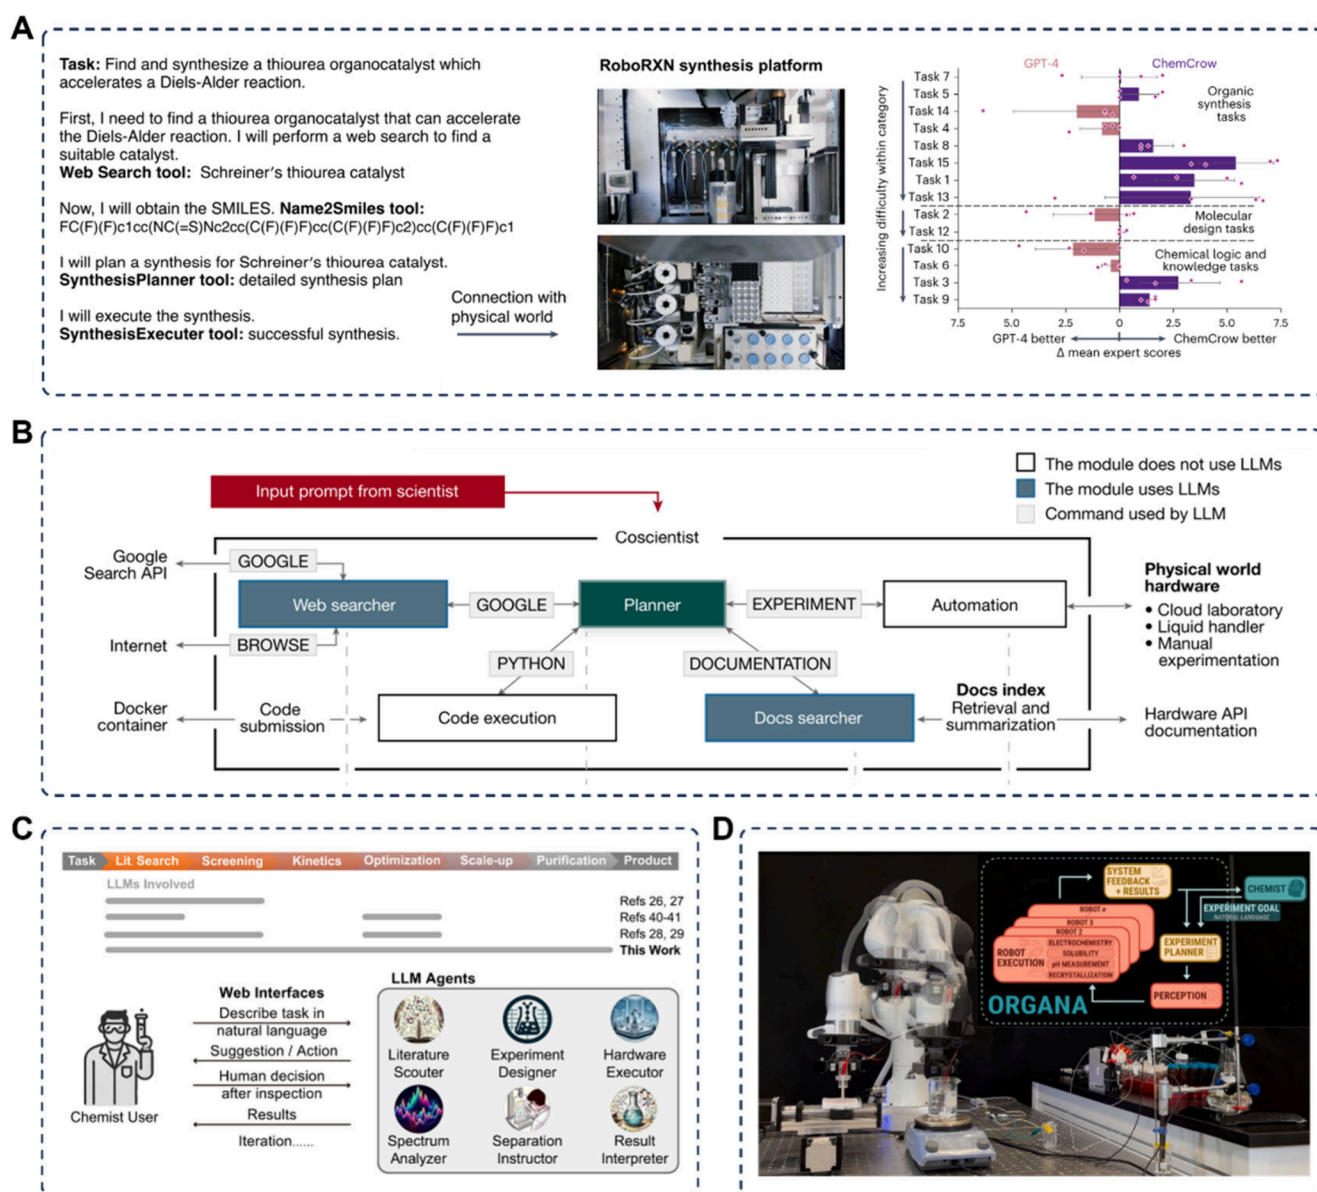


Figure 11. Agentic workflows in materials science. (A) ChemCrow coordinates specialized software tools and the RoboRXN platform to autonomously plan and execute syntheses, outperforming unaugmented LLMs. Adapted from ref 30. Copyright 2024 under CC BY 4.0 license by the Authors, Springer Nature. (B) Coscientist integrates documentation search with code execution to autonomously generate Python scripts that control laboratory instruments. Adapted from ref 28. Copyright 2023 under CC BY 4.0 license by the Authors, Springer Nature. (C) Multiagent orchestration. The LLM-RDF framework distributes research tasks across specialized agents (e.g., Designer, Executor) for collaborative experiment planning and execution. Adapted with permission from ref 143. Copyright 2024 Springer Nature. (D) Embodied intelligence ORGANA combines an LLM planner with visual perception and robotic motion control to manage parallelized physical workflows. Adapted with permission from ref 146. Copyright 2025 Elsevier.

ces robustness to specialized notation; (3) schema-based extraction creates linkable structured knowledge; and (4) standardized ontologies enable execution and reproducibility. Remaining challenges include coverage gaps in under-represented subfields, the risk that text-derived representations may systematically reinforce existing research biases rather than facilitate paradigm-breaking discovery,¹⁴² errors from long context or multimodal inputs, and lack of uncertainty estimation. Addressing these issues will require hybrid systems that couple statistical representations to explicit schemas and tool-aware reasoning. Ultimately, the representational infrastructure outlined here forms the foundation for the agentic AI systems discussed in section 3.2, systems that not only

understand scientific language but can act upon it to reason, plan, and discover.

3.2. Agentic Workflow

The progression from knowledge representation to autonomous scientific action marks a fundamental transition in the evolution of AI for materials science. While Section 3.1 discussed how scientific information can be encoded and structured in machine-interpretable forms, this section explores how such knowledge becomes actionable. Here, artificial intelligence systems cease to be passive repositories of understanding and begin to function as agents—entities capable of reasoning about experimental goals, formulating executable plans, and carrying out physical or computational

tasks. These agentic systems constitute the executive layer of scientific AI, converting abstract representations into verifiable material outcomes and completing the conceptual loop from “knowing” to “doing.”

At the heart of this transformation lies the concept of plan binding, the process of mapping an AI-generated scientific plan, typically expressed in natural language or symbolic reasoning steps, onto executable and verifiable actions in physical or computational environments. In materials science, this binding process translates a linguistic instruction such as “mix precursor A and B, anneal at 800 °C, then characterize the phase” into structured, machine-readable commands that directly control laboratory hardware (e.g., through χ DL or Opentrons API) or simulation workflows. By linking semantic intent to deterministic actions, plan binding ensures that every AI-generated decision can be executed, logged, and validated—bridging the gap between linguistic reasoning and material manipulation.

Tool-augmented architectures operationalize this concept by coupling large language models (LLMs) with domain-specific tools for computation, retrieval, and control. **ChemCrow** embodies a *ReAct-style* reasoning loop (“Thought → Action → Observation”), where a GPT-4 backbone coordinates 18 specialized chemistry tools including molecular analysis, literature search, and robotic synthesis via RoboRXN (Figure 11A).³⁰ This design allowed the system to autonomously plan and execute the synthesis of DEET and three organocatalysts, while coordinating with a human-in-the-loop evaluator. Similarly, **Coscientist** integrates GPT-4 with a Python interpreter, web API search, and hardware documentation database to conduct and troubleshoot real experiments. When an initial code failed to operate a heater–shaker module, the agent retrieved the relevant API documentation through embedding search, corrected its code, and successfully executed a cross-coupling reaction (Figure 11B).²⁸ Scientific reasoning thus becomes actionable only when linguistic intent is *bound* to verifiable interfaces.

Beyond single-agent reasoning, increasing research complexity has driven the evolution of orchestrated multiagent systems capable of managing entire workflows. The LLM-RDF framework employs six specialized GPT-4 agents—Literature Scouter, Experiment Designer, Hardware Executor, Spectrum Analyzer, Separation Instructor, and Result Interpreter—to autonomously plan and perform a Cu/TEMPO-catalyzed alcohol oxidation reaction (Figure 11C).¹⁴³ By coupling Bayesian optimization with robotic execution on an Opentrons OT-2 and Big Kahuna platform, it demonstrated how distributed cognition can achieve closed-loop decision-making across experiment, analysis, and optimization. In contrast, the **AI-Chemist** represents a monolithic but deeply integrated system uniting natural-language planning, machine-reading of literature, molecular simulation, and mobile robotics.^{144,145} Operating across 14 workstations for synthesis, characterization, and performance testing, the **AI-Chemist** autonomously discovered an optimized high-entropy-alloy catalyst for the oxygen-evolution reaction, iteratively updating DFT-based predictions with experimental feedback through Bayesian optimization. The same orchestration principle underlies systems such as **ORGANA**, which integrates an LLM planner, visual perception, and robotic motion control to execute parallelized workflows for electrochemical characterization. By solving task-and-motion planning and scheduling simultaneously, **ORGANA** reduced total experimental time by 21%,

underscoring how agentic orchestration enhances both autonomy and efficiency (Figure 11D).¹⁴⁶

These architectures collectively span a spectrum of autonomy determined by the degree of human involvement and the coupling between reasoning and execution. At one end are *human-in-the-loop copilots* that augment rather than replace scientific judgment. The GPT-4 Reticular Chemist guides researchers through iterative MOF synthesis, updating its planning strategy based on human feedback after each experimental cycle.²⁷ **CREST** enables voice-driven control of robotic instruments, allowing experimentalists to operate complex workflows via natural language without coding.¹⁴⁷ This agentic interface paradigm has since extended to characterization instrumentation, with LLM-driven systems enabling conversational control of scanning probe microscopes and voice-controlled experiments at synchrotron beamlines.^{148,149} At the other end are *closed-loop systems* capable of sustained independent operation. The **Artificial Chemist**¹⁵⁰ autonomously optimized quantum-dot synthesis using in situ optical feedback and Bayesian optimization, completing 11 target compositions within 30 h, whereas the **A-Lab**⁸⁵ executed 355 solid-state synthesis experiments over 17 days, successfully producing 71% of its 58 targets. By iteratively learning from failed experiments, the **A-Lab** leveraged precomputed thermodynamic data and active learning to optimize synthesis routes, demonstrating adaptive experimental reasoning on scale. Between these two poles lie hybrid systems such as **MatPilot**,¹⁵¹ which blends human conceptual oversight with AI-driven automation, establishing an early prototype for collaborative scientific reasoning.

Recent work has begun to formalize and expand this middle ground, moving beyond ad hoc human oversight toward structured frameworks in which human expertise actively shapes the AI’s exploration strategy. Biswas et al. introduced the Bayesian Optimized Active Recommender System (**BOARS**).¹⁵² Harris et al. further formalized this supervisory role through a dual-Gaussian-process architecture.⁷⁴ At the instrument level, the **SEEK** framework demonstrated how domain expertise can be formally encoded into active learning loops.¹⁵³ Collectively, these developments indicate that the appropriate level of human oversight in autonomous experimentation is not a fixed threshold but a variable calibrated to the risk, reversibility, and novelty of each experimental decision—from post hoc review for routine high-throughput screening to active human-in-the-loop control for safety-critical or irreversible transformations—a principle elaborated further in Section 4.4.

As AI agents gain autonomy, the demand for reflexivity—the ability to detect, explain, and correct their own errors—becomes critical. Several frameworks now embed self-evaluation and repair mechanisms to enhance reproducibility and safety. **AutoBA** automates multiomics analysis with an automatic-code-repair loop that captures runtime errors and regenerates code until successful execution, raising its end-to-end success rate from 65 to 87.5%.¹⁵⁴ In symbolic reasoning, **Baldur** employs a dual-LLM “generate-and-repair” workflow, where one model proposes a formal proof and another revises failed attempts using diagnostic feedback.¹⁵⁵ Laboratory-scale reflexivity is exemplified by **BioPlanner**,¹⁵⁶ which formalizes experiment planning by translating natural-language protocols into a closed set of pseudofunctions (e.g., `add_and_mix()`, `spin_sample()`), enabling automatic verification of plan correctness. Such formalization directly parallels the plan

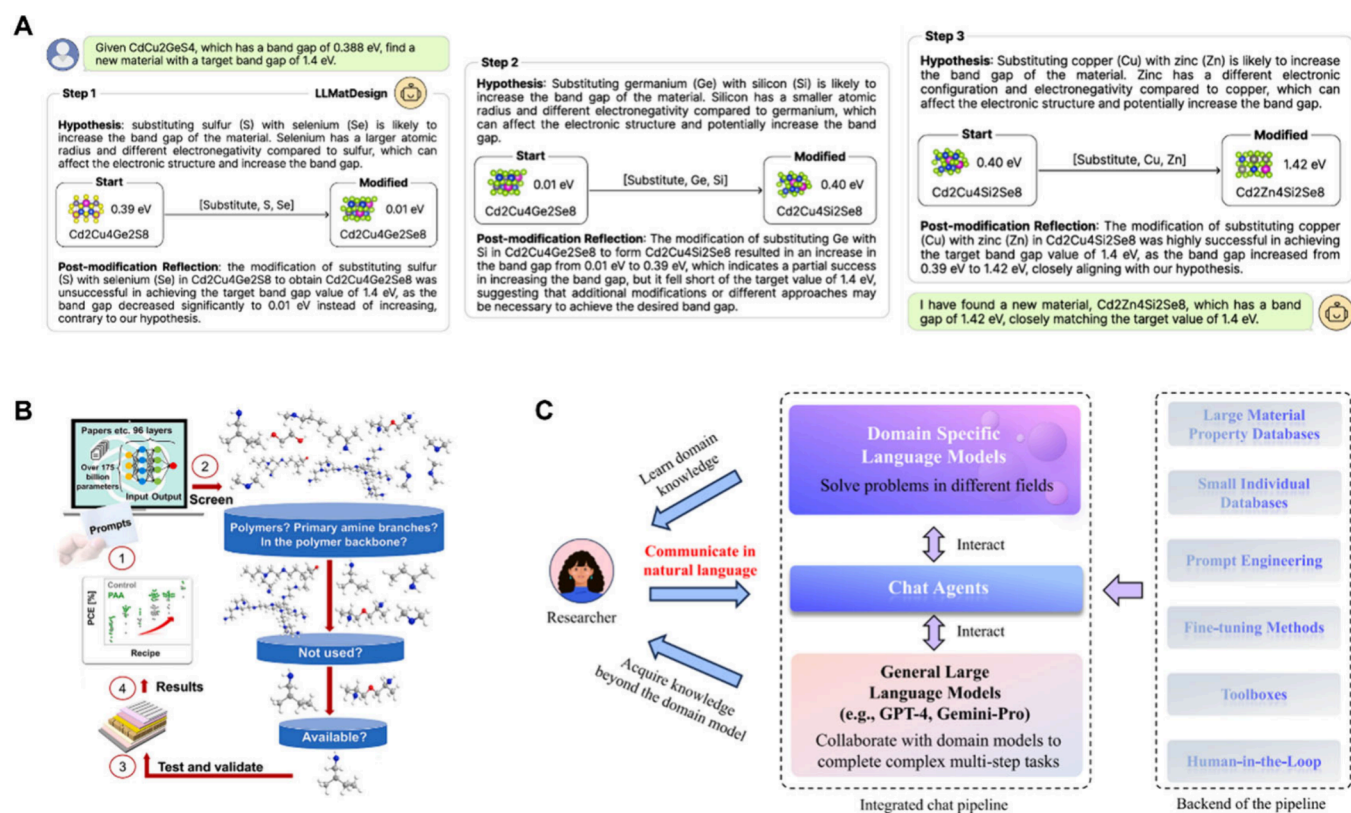


Figure 12. Human-AI collaborative reasoning. (A) Self-reflective diagnostic reasoning. The model mimics the scientific method by formulating explicit hypotheses for band gap tuning, reflecting on computational failures, and self-correcting its strategy to achieve the target property. Adapted with permission from ref 159. Copyright 2024 the Authors. (B) Hypothesis generation via analogical reasoning. Adapted from ref 160. Copyright 2024 under CC BY 4.0 license by the Authors, Elsevier. (C) Context-aware justification. The system bridges qualitative human intent with quantitative design, providing reasoning-based explanations for why a specific material is recommended for a real-world application. Adapted from ref 161. Copyright 2025 under CC BY 3.0 license by the Authors, The Royal Society of Chemistry.

binding concept, ensuring that agentic reasoning yields auditable, deterministic outcomes rather than free-form text.

Taken together, these developments mark the emergence of a new operational layer in scientific AI—systems that do not merely predict or analyze but actively participate in the experimental cycle. By coupling linguistic reasoning with structured, executable action through plan binding and tool augmentation, AI agents transform encoded knowledge into verifiable experimentation. Their capacity to plan, act, reflect, and adapt foreshadows a new paradigm of materials research where discovery becomes both autonomous and explainable. Yet, true scientific value arises not merely from automation but from *understanding*—a transition explored in the next section (3.3), which focuses on coreasoning as the foundation for trustworthy AI in science.

3.3. Human–AI Collaborative Reasoning

Human–AI collaboration in materials science spans a broad continuum. Where previous sections addressed knowledge encoding (Section 3.1) and the orchestration of autonomous experimental workflows with human oversight at the execution level (Section 3.2), this section turns to reasoning-level collaboration—the discursive and interpretive layer in which language models serve not as automated actors but as cognitive partners, helping scientists reason about mechanisms, reconcile conflicting evidence, and generate new hypotheses. This mode of interaction leverages the intrinsic strengths of large language models (LLMs): synthesizing heterogeneous information, constructing analogies, articulating conceptual explanations,

and engaging in iterative dialogue. Unlike the “hard interpretability” of Section 2.4, which aims to uncover quantitative laws or mechanistic descriptors, collaborative reasoning represents a softer epistemic process grounded in narrative understanding, conceptual framing, and hypothesis exploration.

A primary manifestation of this capability arises from domain-adapted scientific language models. By leveraging techniques such as fine-tuning on scientific corpora or Retrieval-Augmented Generation (RAG) to ground their knowledge, these systems evolve from static information retrievers into reasoning engines capable of articulating the logic behind their decisions.^{157,158} A quintessential example of this explanatory capability is the LLMatDesign framework.¹⁵⁹ Unlike traditional black-box optimization algorithms, this system is designed to mimic human scientific reasoning by formulating an explicit chemical hypothesis before proposing structural modifications. For instance, when tasked with tuning the band gap of a semiconductor (CdCu₂GeS₄), the model did not simply guess parameters; it proposed substituting sulfur (S) with selenium (Se), explicitly reasoning that selenium’s “larger atomic radius and different electronegativity” would alter the electronic structure to change the band gap. Crucially, these systems demonstrate diagnostic reasoning through self-reflection loops. When the initial substitution failed to achieve the target property (the band gap decreased instead of increased), the model analyzed the discrepancy between its hypothesis and the computational result. It acknowledged the

failure, adjusted its chemical intuition, and formulated a new strategy (substituting germanium with silicon and copper with zinc), which eventually succeeded. By recognizing periodic trends and causally explaining its iterative decision-making process, such models provide transparent rationales that allow researchers to validate the scientific soundness of AI-generated concepts (Figure 12A).

Beyond diagnostic analysis, collaborative reasoning plays a pivotal role in the ideation and hypothesis generation phase of discovery. Here, the LLM functions as a brainstorming partner, capable of transferring mechanistic logic from one chemical system to another via analogical reasoning. A striking demonstration of this capacity is the work by Chen et al., who utilized ChatGPT to identify novel surface passivators for perovskite solar cells. Rather than simply retrieving known materials, the researchers engaged the AI in an iterative dialogue to search for untested candidates that shared specific functional traits (primary amines) with established passivators but offered distinct structural advantages (polymeric backbones). The model reasoned that polyallylamine (PAA)—a molecule never previously used for this specific purpose—would effectively reduce surface recombination due to its structural similarity to known passivators and high density of amine groups. This AI-generated hypothesis was subsequently validated experimentally, yielding a measurable increase in device efficiency. This case exemplifies how conversational AI can bridge the gap between “known physics” and “novel application” by speculating on structure–property relationships in a way that parallels human intuition (Figure 12B).¹⁶⁰

Furthermore, collaborative reasoning enables context-aware inverse design, where LLMs translate qualitative human requirements into quantitative material constraints. Liu and Wen illustrate this with ElaTBot, a domain-specific model designed for elastic constant tensor prediction. Unlike standard regression models that output mere numbers, ElaTBot allows researchers to articulate complex, multiobjective design goals in natural language—such as seeking an orthopedic material that matches the bulk modulus of human bone while maintaining specific biocompatibility standards (e.g., LD50 limits). The system processes these constraints to recommend specific alloy compositions (e.g., Y₃Al) and, crucially, provides a textual rationale explaining *why* the material fits the criteria based on its corrosion resistance and bulk modulus. This transition from “predicting properties” to “justifying recommendations” marks a fundamental shift: the AI does not merely optimize a target function but actively reasons about the material’s suitability for a real-world application context (Figure 12C).¹⁶¹

Taken together, human–AI collaborative reasoning completes the arc of scientific intelligence developed throughout this section. While foundation models (Section 3.1) encode the vast topology of scientific knowledge and agentic systems (Section 3.2) operationalize it through autonomous execution, collaborative reasoning provides the essential interpretive layer. Whether diagnosing a failed DFT calculation, identifying a serendipitous chemical analogy for surface passivation, or justifying a material choice for biomedical implants, these systems augment the *epistemic* processes of science. They do not merely automate labor; they expand the conceptual search space available to the researcher. As these domain-adapted models become increasingly capable of articulating the “why” behind their predictions, the boundary between human intuition and machine logic blurs, pointing toward a future

of genuinely shared scientific cognition where AI acts as a coauthor of the scientific hypothesis itself. Yet current benchmarks counsel realism: MaScQA shows that state-of-the-art LLMs reach only undergraduate-level competence on materials science questions, faltering on the conceptual nuances that domain experts navigate routinely,¹³⁹ and SUPERChem confirms that multimodal models still fall short of human experts in higher-order chemical reasoning.¹⁶² The collaborative paradigm showcased above is therefore not merely a design preference but a practical necessity at the present stage of the technology.

4. OUTLOOK AND PERSPECTIVES

The preceding sections have charted the evolution of AI4Mat along two complementary trajectories—highlighting both transformative advances and the specific limitations that temper current progress at each stage. While these technologies have successfully accelerated individual stages of discovery, the ultimate vision of a self-driving, closed-loop research paradigm remains a systemic challenge. Moving forward, we identify five critical dimensions that will shape the future trajectory of AI4Mat.

4.1. Data Ecosystems and Automated Infrastructure

The efficacy of AI in materials science is intrinsically bound by the quality, quantity, and accessibility of data. While the AI boom has been propelled by the open-source ethos of the computer science community, exemplified by shared repositories like GitHub and Hugging Face, materials science remains largely constrained by data silos and nonstandardized reporting. To unlock the full potential of AI, the research paradigm must evolve from artisanal data collection to an industrialized intelligent ecosystem. This transformation necessitates a tripartite strategy: standardization, automated generation, and an algorithmic adaptation to scarcity.

Standardization and Knowledge Engineering: The first imperative is to dismantle data silos through community-wide standardization. As discussed in Section 3.1, the heterogeneity of experimental data, spanning diverse characterization modalities and unstandardized text descriptions—poses a major barrier to machine learning. We must advocate for the universal adoption of structured ontologies and description languages, such as the Chemical Description Language (XDL)¹³⁶ and schema-based extraction pipelines¹³⁰ reviewed earlier. Furthermore, this standardization must extend beyond successful results to include failed experiments and negative data.³¹ In current literature, the omission of failed syntheses introduces a severe survivorship bias; systematically capturing this “dark data” is essential for training models that can discern the boundaries of chemical feasibility. AI itself plays a critical role here as a “data engineer”: utilizing the Knowledge Engineering tools (Section 3.1) to automatically normalize, align, and clean multimodal data from disparate sources, thereby converting raw experimental noise into machine-actionable assets.

Industrialized Data Generation via Automated Laboratories: Standardization alone cannot solve the data volume bottleneck. The field must transition toward decentralized, high-throughput automated laboratories. The success of platforms like the A-Lab⁸⁵ and RoboChem⁷⁵ (analyzed in Section 2.2.2) demonstrates that integrating robotics with active learning agents can transform data generation from a manual bottleneck into an industrialized pipeline. These “self-driving labs” do not

merely accelerate experimentation; they ensure epistemic consistency by reducing human variability and generating high-fidelity, standardized data sets at a scale previously unattainable. This shift envisions a future of “cloud laboratories”, analogous to cloud computing, where physical experimentation is decoupled from the researcher’s location. By enabling scientists to submit experimental protocols via code to centralized automated facilities, this democratizes access to state-of-the-art infrastructure, lowering the entry barrier for resource-limited groups. Furthermore, it ensures that every data point is generated under strictly monitored conditions, replacing ambiguous manual reporting with immutable, machine-readable provenance logs.

Despite the impressive demonstrations reviewed in Section 2.2, the broader adoption of autonomous experimentation in typical academic laboratories remains hindered by several interconnected bottlenecks. First, hardware heterogeneity poses a persistent integration challenge: instruments from different vendors lack standardized communication interfaces, and the engineering effort required to integrate even two or three devices into a closed loop often exceeds the expertise and resources available in a single research group. Second, the interdisciplinary talent gap is acute—effective deployment requires simultaneous expertise in machine learning, hardware automation, and domain-specific materials science, a combination rarely found in a single individual and difficult to assemble in small academic teams. Third, the economic barrier is substantial: the capital cost of building an autonomous platform, including robotics, sensors, and computing infrastructure, remains prohibitive for most academic budgets, and unlike shared instruments such as electron microscopes, autonomous laboratories are typically designed for specific workflows and are difficult to repurpose. Software-level solutions, such as cross-platform abstraction layers that encapsulate heterogeneous hardware behind unified APIs (see Section 2.2.2), offer a promising pathway to lower these barriers by decoupling algorithmic development from hardware specifics.

Algorithmic Resilience in Data-Scarce Regimes: Concurrently, we must recognize that “Big Data” will remain elusive for many specialized or nascent material systems due to the immense cost of physical validation. In the broader AI domain, foundation models are typically trained on trillions of tokens or billions of images. Even within computational materials science, simulation databases like GNoME or the Open Catalyst Project can reach the order of 10^6 to 10^8 data points. In stark contrast, experimental data sets are constrained by the physical laws of synthesis and characterization time. A typical manual study might yield fewer than 100 data points (10^2), and even state-of-the-art autonomous laboratories (e.g., A-Lab) generate data on the scale of 10^2 to 10^3 samples per campaign. Given this persistent magnitude gap, expecting purely data-driven models to solve every problem is unrealistic. A parallel strategic priority is the development of sample-efficient algorithms tailored for data-scarce regimes. Paradigms like few-shot learning and meta-learning offer theoretical avenues to generalize from minimal examples, their rigorous adaptation to the complexities of materials science remains a critical frontier. Rather than viewing data scarcity as a temporary bottleneck to be solved by volume, the community must prioritize the engineering of “Small Data” intelligence, algorithms capable of leveraging physical priors to extract maximum entropy from sparse data sets. This methodological

evolution is essential to ensure that AI serves not only data-rich domains but also the “small data” reality of cutting-edge experimental discovery. Compounding the data volume challenge is a less recognized but equally critical issue: cross-platform transferability. The majority of AI models in materials characterization have been developed and validated within single instrumental setups, where instrument-specific variations in calibration, detector response, sample preparation, and environmental conditions can introduce systematic distribution shifts. While domain-specific pretraining—such as the MicroNet repository for microscopy (Section 2.3.1)—has demonstrated that aligned feature representations can dramatically reduce data requirements, the community currently lacks the large-scale, multi-instrument benchmark data sets that would enable systematic evaluation of cross-platform generalization. The development of such benchmarks, analogous to the multisite clinical imaging data sets that catalyzed progress in medical AI, represents an important data infrastructure priority.

4.2. Balancing Generality and Specialization

The current trajectory of AI is dominated by the pursuit of AGI, driven by the dogma that “the larger the model, the better the performance.” While this scaling law holds for natural language and broad cognitive tasks, applying it indiscriminately to materials science is premature and fraught with epistemic risks. Unlike creative writing, scientific inquiry demands strict exactitude; a hallucinated chemical structure or an erroneous band gap prediction is not a stylistic variant but a factual failure. Therefore, the near-future of AI4Mat should not rely on a monolithic generalist model, but rather on a symbiotic ecosystem that balances the breadth of agents with the depth of experts.

This architecture necessitates a hierarchical division of labor that mirrors the biological metaphor illustrated in Figure 2. At the foundational level, Task-Specific AI functions as the specialized “Organs” of the system. Just as distinct organs perform vital biological functions with high specificity, specialized models reviewed in Section 2, provide the rigorous, grounded “truth” required for scientific tasks. These models consistently outperform generalist counterparts in precision and data efficiency within their defined domains. Conversely, Generalist AI (e.g., Large Language Models) acts as the integrative “Blood”—or the “Orchestrator”—that connects these isolated functional modules. Its primary utility in science is not to solve Schrödinger’s equation directly, but to manage the flow of reasoning and information. As discussed in Section 3, the generalist agent parses natural language intent, plans multistep workflows, and routes tasks to the appropriate specialized tools. Consequently, the paradigm shift in AI4Mat lies not in replacing the specialized AI with a generalist AI, but in optimizing the circulation between them—ensuring that the reasoning capabilities of agents are effectively grounded in the quantitative reliability of expert models.

4.3. Bridging Simulation, Experiment, and Theories

A preponderance of current AI4Mat research is concentrated on predictive modeling and virtual screening (as reviewed in Section 2.1), primarily because simulation data is structured, controllable, and governed by known equations. Landmark foundation models, such as GNoME⁵¹ and MatterGen,¹⁶³ have demonstrated the immense power of deep learning by expanding the known crystal space to millions of potential structures. However, there is a big gap between simulations

and real-world experiments. Simulations typically rely on idealized assumptions: perfect crystal lattices, stoichiometric purity, and thermodynamic equilibrium. In contrast, real-world synthesis is inherently stochastic, dominated by nonequilibrium dynamics, defect structures, and the complex kinetics of precursor decomposition. Therefore, to realize the full value of these large-scale predictions, the field must address the “Valley of Death” that separates digital discovery from physical realization.

Embodied Intelligence¹⁶⁴ offers a promising pathway to bridge this chasm. Rather than relying solely on precomputed data sets, integrating AI agents with robotic hardware and sensory feedback systems^{73,85} introduces the possibility of learning directly from physical interaction. Unlike simulations that operate in idealized environments, such embodied agents have the potential to navigate the “dirty” data of reality, managing noise, impurities, and equipment variability. By attempting to close the loop between digital prediction and physical execution, this approach aims to capture the implicit craft of synthesis, exploring whether AI can evolve from a tool that predicts theoretical stability into an agent that masters experimental realizability.

Bridging the physical gap is only half the scientific mission. Materials science is fundamentally distinct from pure manufacturing; its ultimate objective is not merely utility (making things work), but understanding (knowing why they work). An AI agent that successfully navigates the “Valley of Death” through opaque heuristics might serve as a “skilled artisan”, but it fails as a scientific partner if it cannot articulate the principles behind its success. Scientific interpretability requires the extraction of structured domain knowledge, explicit physical laws and causal mechanisms.

We envision AI systems that evolve from Oracles (that provide probabilistic answers) to Agents (that execute synthesis) and Theorists (that build structured knowledge). AI should assist in distilling the complex, “dirty” data of reality into parsimonious physical laws. This ensures that the AI revolution yields not just an inventory of new substances, but an expansion of the fundamental principles governing matter.

Yet understanding alone does not close the final gap. Cross-scale fabrication—translating atomic-level design into macroscopic function—remains largely artisanal, each transfer step managed through ad hoc optimization rather than unified methodology. To move beyond this, we propose Material Compilation Science and Engineering (MCSE): a disciplinary vision in which fabrication across scales is organized into modular, compilable stages with certain rules, enabling macroscopic function to reliably emerge from microscopic design. Realizing MCSE will demand the convergence of predictive modeling, autonomous experimentation, and theory extraction into a coherent engineering discipline—the point at which AI4Mat matures from accelerating discovery to mastering fabrication.

4.4. Collaborative Ecosystems in the Age of AI4Mat

From a macroscopic perspective, the evolution of AI4Mat extends beyond specific technical breakthroughs to the construction of collaborative ecosystems. Addressing multifaceted research problems demands synergy at three levels. First, Human-to-Human collaboration must evolve beyond simple teamwork to the cultivation of interdisciplinary talent. Bridging the gap between chemistry, physics, and computer science requires researchers who possess both deep domain

expertise and technical fluency. Second, AI-to-AI collaboration is materializing through multiagent systems. By coordinating specialized agents—each dedicated to distinct tasks such as literature mining, code generation, or property prediction—these systems can tackle complex problems that are insoluble by single models. Third, Human-AI collaboration is bidirectional: on one hand, AI tools undertake labor-intensive execution tasks, liberating scientists from routine operations to focus on creativity. On the other hand, human intuition remains superior in frontier domains; for instance, the recent SUPERChem¹⁶² benchmark demonstrates that even top-tier models lag behind human experts in high-order reasoning tasks. Expert knowledge is thus indispensable, not only for verifying AI outputs but for providing the high-quality feedback required to continuously train and align models with physical reality.

Furthermore, as we empower AI with agentic capabilities to manipulate physical matter, humans must serve as the critical ethical and safety guardrails. The potential for “dual-use” risks requires rigorous human oversight to define the boundaries of safe exploration. Ultimately, the future of materials science is not an automated loop that runs without humans, but a dynamic ecosystem where collaborative intelligence drives innovation within meaningful and ethical frameworks.

As autonomous experimentation platforms advance from proof-of-concept demonstrations to routine deployment, the question of appropriate human oversight becomes increasingly urgent. Recent work on human-supervised autonomous experimentation—including the dynamic optimization, active oversight, and expert-knowledge-informed frameworks discussed in Section 3.2—provides initial models for graduated autonomy, where the level of human supervision is calibrated to the risk and novelty of each experimental decision. Establishing community guidelines for such tiered supervision, particularly for experiments involving hazardous materials, irreversible transformations, or expensive instrumentation, represents an important near-term priority.

4.5. From AI4Mat to Mat4AI

The preceding sections extensively explored how AI transforms the research paradigm of materials science (AI4Mat). Conversely, material innovation is fundamentally critical for the evolution of AI (Mat4AI), establishing a bidirectional symbiosis between the two fields.

On the one hand, the explosive growth of AI models is currently outpacing the capabilities of underlying hardware, with traditional CMOS scaling approaching its physical limits.^{165,166} Here, materials science serves as the physical bedrock for the next generation of computing. The discovery of novel low-dimensional materials,^{167–169} carbon nanotubes,^{170–172} and 2D semiconductors^{173,174} offers the potential for transistors with superior speed and energy efficiency, while advancements in memristive and phase-change materials are critical for realizing neuromorphic architectures.¹⁷⁵ Indeed, several of these material platforms have already progressed from device-level demonstrations to functional AI hardware: carbon nanotube field-effect transistors have been integrated into a tensor processing unit capable of convolutional neural network inference,¹⁷⁶ while single-chip silicon nitride photonic processors now achieve subnanosecond inference latency on a commercial CMOS foundry,¹⁷⁷ and memristive crossbar arrays¹⁷⁸ and halide perovskite neuromorphic devices have

demonstrated hardware-based neural network inference, including transformer-class architectures.¹⁷⁹

Simultaneously, for embodied intelligence, advanced materials are indispensable for constructing the physical agents themselves. High-performance composite materials and soft actuators provide the essential structural integrity and flexibility required for robotics and autonomous systems.^{180–182} Furthermore, the sustainability of AI relies heavily on energy materials. Addressing the massive energy consumption of AI requires breakthroughs not just in storage,^{183,184} but fundamentally in energy generation. From high-efficiency photovoltaics to extreme-environment materials for nuclear fusion, materials innovation is the primary bottleneck preventing the limitless, clean power supply needed to satisfy the voracious appetite of future AI models.^{185–187} By developing these advanced material systems, our field does not merely utilize AI; it constructs the essential “body” required for future AI generations to function sustainably.

However, the contribution of our discipline extends beyond this physical substrate to the very core of cognitive evolution. Unlike the noisy and often hallucinated data of the open web, materials science deals with the authentic behavior of matter, providing rigorous, high-fidelity data sets governed by physical laws. These unique characteristics position materials science as the vanguard of a broader evolution in artificial intelligence. The development of AI has completed its “first half” of learning from the Internet’s stock knowledge and is now entering the “second half” of creating incremental knowledge through interaction with the physical world. Since scientific research, epitomized by materials science, represents humanity’s most efficient and high-quality process for knowledge generation, the convergence of AI and science represents not merely a frontier of application, but a key driving force for upgrading AI’s core capabilities.

■ AUTHOR INFORMATION

Corresponding Authors

Liu Qian – School of Materials Science and Engineering, Peking University, Beijing 100871, China; Beijing Science and Engineering Center for Nanocarbons, Beijing National Laboratory for Molecular Sciences, College of Chemistry and Molecular Engineering, Peking University, Beijing 100871, China; Email: qianliu-cnc@pku.edu.cn

Jin Zhang – School of Materials Science and Engineering, Peking University, Beijing 100871, China; Beijing Science and Engineering Center for Nanocarbons, Beijing National Laboratory for Molecular Sciences, College of Chemistry and Molecular Engineering, Peking University, Beijing 100871, China; orcid.org/0000-0003-3731-8859; Email: jinzhang@pku.edu.cn

Authors

Yue Li – School of Materials Science and Engineering, Peking University, Beijing 100871, China; Beijing Science and Engineering Center for Nanocarbons, Beijing National Laboratory for Molecular Sciences, College of Chemistry and Molecular Engineering and School of Physics, Peking University, Beijing 100871, China; orcid.org/0000-0002-3972-0692

Shurui Wang – Beijing Science and Engineering Center for Nanocarbons, Beijing National Laboratory for Molecular Sciences, College of Chemistry and Molecular Engineering, Peking University, Beijing 100871, China

Jianping Wang – Beijing Science and Engineering Center for Nanocarbons, Beijing National Laboratory for Molecular Sciences, College of Chemistry and Molecular Engineering, Peking University, Beijing 100871, China; orcid.org/0009-0009-5663-9157

Complete contact information is available at:
<https://pubs.acs.org/10.1021/acs.chemrev.6c00012>

Notes

The authors declare no competing financial interest.

Biographies

Yue Li received his B.S. (2020) and Ph.D. (2025) from the School of Physics, Peking University. He conducted his doctoral research under the supervision of Prof. Jin Zhang and Prof. Ziqiang Zhao, dedicated to integrating machine learning and automation into the controlled synthesis of carbon nanotubes. He is currently a Postdoctoral Research Fellow at Nanyang Technological University, in the group of Prof. Zheng Liu. His research interests focus on AI for Science, particularly the application of AI methods in the synthesis and engineering of low-dimensional materials.

Shurui Wang received her B.S. from Peking University in 2024, where she is currently pursuing her Ph.D. in Physical Chemistry under the supervision of Prof. Jin Zhang and Dr. Liu Qian. Her current research focuses on the AI-guided catalyst design and scalable synthesis of carbon nanotube arrays.

Jianping Wang received his B.S. in Chemical Physics from the University of Science and Technology of China (2024). He is currently pursuing his Ph.D. in Physical Chemistry at Peking University under the supervision of Prof. Jin Zhang and Dr. Liu Qian. His research focuses on leveraging AI for the structure-controlled growth of carbon nanotubes.

Liu Qian is an assistant researcher at the School of Materials Science and Engineering, Peking University. She has long been dedicated to the controlled preparation of horizontally aligned carbon nanotube arrays and deeply explores the convergence of materials science and AI technology. She has achieved breakthroughs in the structural-controlled growth of carbon nanotubes, wafer-scale high-density carbon nanotube arrays, and AI-enabled precise synthesis of carbon nanotubes.

Jin Zhang is an Academician of the Chinese Academy of Sciences, a Fellow of the Royal Society of Chemistry, and a Boya Chair Professor at Peking University. His research primarily focuses on the controlled synthesis and spectroscopic characterization of carbon nanomaterials, with notable contributions to structure-controlled growth mechanisms, electronic-grade carbon nanotube arrays and high-performance carbonene fibers. Recently, he is leading the strategic integration of artificial intelligence with materials science, aiming to establish new paradigms for AI-driven scientific discovery.

■ ACKNOWLEDGMENTS

This work was supported by the National Natural Science Foundation of China (52021006 and 22494641), the Beijing National Laboratory for Molecular Sciences (BNLMS-CXTD-202001) and the Shenzhen Science and Technology Innovation Commission (KQTD20221101115627004).

■ REFERENCES

(1) Kirkpatrick, P.; Ellis, C. Chemical Space. *Nature* **2004**, *432* (7019), 823–823.

- (2) Bohacek, R. S.; McMartin, C.; Guida, W. C. The Art and Practice of Structure-Based Drug Design: A Molecular Modeling Perspective. *Med. Res. Rev.* **1996**, *16* (1), 3–50.
- (3) Xu, Y.; Liu, X.; Cao, X.; Huang, C.; Liu, E.; Qian, S.; Liu, X.; Wu, Y.; Dong, F.; Qiu, C.-W.; Qiu, J.; Hua, K.; Su, W.; Wu, J.; Xu, H.; Han, Y.; Fu, C.; Yin, Z.; Liu, M.; Roepman, R.; Dietmann, S.; Virta, M.; Kengara, F.; Zhang, Z.; Zhang, L.; Zhao, T.; Dai, J.; Yang, J.; Lan, L.; Luo, M.; Liu, Z.; An, T.; Zhang, B.; He, X.; Cong, S.; Liu, X.; Zhang, W.; Lewis, J. P.; Tiedje, J. M.; Wang, Q.; An, Z.; Wang, F.; Zhang, L.; Huang, T.; Lu, C.; Cai, Z.; Wang, F.; Zhang, J. Artificial Intelligence: A Powerful Paradigm for Scientific Research. *Innovation* **2021**, *2* (4), 100179.
- (4) Wang, H.; Fu, T.; Du, Y.; Gao, W.; Huang, K.; Liu, Z.; Chandak, P.; Liu, S.; Van Katwyk, P.; Deac, A.; Anandkumar, A.; Bergen, K.; Gomes, C. P.; Ho, S.; Kohli, P.; Lasenby, J.; Leskovec, J.; Liu, T.-Y.; Manrai, A.; Marks, D.; Ramsundar, B.; Song, L.; Sun, J.; Tang, J.; Veličković, P.; Welling, M.; Zhang, L.; Coley, C. W.; Bengio, Y.; Zitnik, M. Scientific Discovery in the Age of Artificial Intelligence. *Nature* **2023**, *620* (7972), 47–60.
- (5) Van Noorden, R.; Perkel, J. M. AI and Science: What 1,600 Researchers Think. *Nature* **2023**, *621* (7980), 672–675.
- (6) Zhou, B.; Ding, N.; Bai, L.; Zhou, H. Advancing AI for Science: From the Revolution of Tools to the Tools for Revolution. *AI Open* **2025**, *6*, 323–328.
- (7) Jumper, J.; Evans, R.; Pritzel, A.; Green, T.; Figurnov, M.; Ronneberger, O.; Tunyasuvunakool, K.; Bates, R.; Židek, A.; Potapenko, A.; Bridgland, A.; Meyer, C.; Kohli, S. A. A.; Ballard, A. J.; Cowie, A.; Romera-Paredes, B.; Nikolov, S.; Jain, R.; Adler, J.; Back, T.; Petersen, S.; Reiman, D.; Clancy, E.; Zielinski, M.; Steinegger, M.; Pacholska, M.; Berghammer, T.; Bodenstein, S.; Silver, D.; Vinyals, O.; Senior, A. W.; Kavukcuoglu, K.; Kohli, P.; Hassabis, D. Highly Accurate Protein Structure Prediction with AlphaFold. *Nature* **2021**, *596* (7873), 583–589.
- (8) 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.; Žemgulyte, A.; Arvaniti, E.; Beattie, C.; Bertolli, O.; Bridgland, A.; Cherepanov, A.; Congreve, M.; Cowen-Rivers, A. I.; Cowie, A.; Figurnov, M.; Fuchs, F. B.; Gladman, H.; Jain, R.; Khan, Y. A.; Low, C. M. R.; Perlin, K.; Potapenko, A.; Savy, P.; Singh, S.; Stecula, A.; Thillaisundaram, A.; Tong, C.; Yakneen, S.; Zhong, E. D.; Zielinski, M.; Židek, A.; Bapst, V.; Kohli, P.; Jaderberg, M.; Hassabis, D.; Jumper, J. M. Accurate Structure Prediction of Biomolecular Interactions with AlphaFold 3. *Nature* **2024**, *630* (8016), 493–500.
- (9) Kates-Harbeck, J.; Svyatkovskiy, A.; Tang, W. Predicting Disruptive Instabilities in Controlled Fusion Plasmas through Deep Learning. *Nature* **2019**, *568* (7753), 526–531.
- (10) Degraeve, J.; Felici, F.; Buchli, J.; Neunert, M.; Tracey, B.; Carpanese, F.; Ewalds, T.; Hafner, R.; Abdolmaleki, A.; De Las Casas, D.; Donner, C.; Fritz, L.; Galperti, C.; Huber, A.; Keeling, J.; Tsimpoukelli, M.; Kay, J.; Merle, A.; Moret, J.-M.; Noury, S.; Pesamosca, F.; Pfau, D.; Sauter, O.; Sommariva, C.; Coda, S.; Duval, B.; Fasoli, A.; Kohli, P.; Kavukcuoglu, K.; Hassabis, D.; Riedmiller, M. Magnetic Control of Tokamak Plasmas through Deep Reinforcement Learning. *Nature* **2022**, *602* (7897), 414–419.
- (11) Seo, J.; Kim, S.; Jalalvand, A.; Conlin, R.; Rothstein, A.; Abbate, J.; Erickson, K.; Wai, J.; Shousha, R.; Kolemen, E. Avoiding Fusion Plasma Tearing Instability with Deep Reinforcement Learning. *Nature* **2024**, *626* (8000), 746–751.
- (12) Lam, R.; Sanchez-Gonzalez, A.; Willson, M.; Wirsberger, P.; Fortunato, M.; Alet, F.; Ravuri, S.; Ewalds, T.; Eaton-Rosen, Z.; Hu, W.; Merose, A.; Hoyer, S.; Holland, G.; Vinyals, O.; Stott, J.; Pritzel, A.; Mohamed, S.; Battaglia, P. Learning Skillful Medium-Range Global Weather Forecasting. *Science* **2023**, *382* (6677), 1416–1421.
- (13) Bi, K.; Xie, L.; Zhang, H.; Chen, X.; Gu, X.; Tian, Q. Accurate Medium-Range Global Weather Forecasting with 3D Neural Networks. *Nature* **2023**, *619* (7970), 533–538.
- (14) Pathak, J.; Subramanian, S.; Harrington, P.; Raja, S.; Chattopadhyay, A.; Mardani, M.; Kurth, T.; Hall, D.; Li, Z.; Azzadnesheli, K.; Hassanzadeh, P.; Kashinath, K.; Anandkumar, A. FourCastNet: A Global Data-Driven High-Resolution Weather Model Using Adaptive Fourier Neural Operators. *arXiv* **2022**, DOI: 10.48550/arXiv.2202.11214 (accessed 2026-04-01).
- (15) Tran, H.; Gurmani, R.; Kim, C.; Pilania, G.; Kwon, H.-K.; Lively, R. P.; Ramprasad, R. Design of Functional and Sustainable Polymers Assisted by Artificial Intelligence. *Nat. Rev. Mater.* **2024**, *9* (12), 866–886.
- (16) Zhang, K.; Gong, X.; Jiang, Y. Machine Learning in Soft Matter: From Simulations to Experiments. *Adv. Funct. Mater.* **2024**, *34* (24), 2315177.
- (17) Du, Y.; Jamasb, A. R.; Guo, J.; Fu, T.; Harris, C.; Wang, Y.; Duan, C.; Liò, P.; Schwaller, P.; Blundell, T. L. Machine Learning-Aided Generative Molecular Design. *Nat. Mach. Intell.* **2024**, *6* (6), 589–604.
- (18) Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; Polosukhin, I. Attention Is All You Need. *arXiv* **2017**, DOI: 10.48550/arXiv.1706.03762.
- (19) Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J. D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; Agarwal, S.; Herbert-Voss, A.; Krueger, G.; Henighan, T.; Child, R.; Ramesh, A.; Ziegler, D.; Wu, J.; Winter, C.; Hesse, C.; Chen, M.; Sigler, E.; Litwin, M.; Gray, S.; Chess, B.; Clark, J.; Berner, C.; McCandlish, S.; Radford, A.; Sutskever, I.; Amodei, D. Language Models Are Few-Shot Learners. In *Advances in Neural Information Processing Systems*; Curran Associates, Inc., 2020; Vol. 33, pp 1877–1901.
- (20) Reed, S.; Zolna, K.; Parisotto, E.; Colmenarejo, S. G.; Novikov, A.; Barth-Maron, G.; Gimenez, M.; Sulsky, Y.; Kay, J.; Springenberg, J. T.; Eccles, T.; Bruce, J.; Razavi, A.; Edwards, A.; Heess, N.; Chen, Y.; Hadsell, R.; Vinyals, O.; Bordbar, M.; Freitas, N. de. A Generalist Agent. *arXiv* **2022**, DOI: 10.48550/arXiv.2205.06175 (accessed 2026-04-01).
- (21) Bubeck, S.; Chandrasekaran, V.; Eldan, R.; Gehrke, J.; Horvitz, E.; Kamar, E.; Lee, P.; Lee, Y. T.; Li, Y.; Lundberg, S.; Nori, H.; Palangi, H.; Ribeiro, M. T.; Zhang, Y. Sparks of Artificial General Intelligence: Early Experiments with GPT-4. *arXiv* **2023**, DOI: 10.48550/arXiv.2303.12712 (accessed 2026-04-01).
- (22) Butler, K. T.; Davies, D. W.; Cartwright, H.; Isayev, O.; Walsh, A. Machine Learning for Molecular and Materials Science. *Nature* **2018**, *559* (7715), 547–555.
- (23) Xie, T.; Grossman, J. C. Crystal Graph Convolutional Neural Networks for an Accurate and Interpretable Prediction of Material Properties. *Phys. Rev. Lett.* **2018**, *120* (14), 145301.
- (24) Schmidt, J.; Marques, M. R. G.; Botti, S.; Marques, M. A. L. Recent Advances and Applications of Machine Learning in Solid-State Materials Science. *Npj Comput. Mater.* **2019**, *5* (1), 83.
- (25) Wiggins, W. F.; Tejani, A. S. On the Opportunities and Risks of Foundation Models for Natural Language Processing in Radiology. *Radiol. Artif. Intell.* **2022**, *4* (4), No. e220119.
- (26) Bommasani, R.; Hudson, D. A.; Adeli, E.; Altman, R.; Arora, S.; von Arx, S.; Bernstein, M. S.; Bohg, J.; Bosselut, A.; Brunskill, E.; et al. On the Opportunities and Risks of Foundation Models. *arXiv* **2021**, DOI: 10.48550/arXiv.2108.07258, August 16, (accessed 2026-04-01).
- (27) Zheng, Z.; Rong, Z.; Rampal, N.; Borgs, C.; Chayes, J. T.; Yaghi, O. M. A GPT-4 Reticular Chemist for Guiding MOF Discovery**. *Angew. Chem., Int. Ed.* **2023**, *62* (46), No. e202311983.
- (28) Boiko, D. A.; MacKnight, R.; Kline, B.; Gomes, G. Autonomous Chemical Research with Large Language Models. *Nature* **2023**, *624* (7992), 570–578.
- (29) Tom, G.; Schmid, S. P.; Baird, S. G.; Cao, Y.; Darvish, K.; Hao, H.; Lo, S.; Pablo-García, S.; Rajaonson, E. M.; Skreta, M.; Yoshikawa, N.; Corapi, S.; Akkoc, G. D.; Strieth-Kalthoff, F.; Seifrid, M.; Aspuru-Guzik, A. Self-Driving Laboratories for Chemistry and Materials Science. *Chem. Rev.* **2024**, *124* (16), 9633–9732.

- (30) Bran, A. M.; Cox, S.; Schilter, O.; Baldassari, C.; White, A. D.; Schwaller, P. Augmenting Large Language Models with Chemistry Tools. *Nat. Mach. Intell.* **2024**, *6* (5), 525–535.
- (31) Raccuglia, P.; Elbert, K. C.; Adler, P. D. F.; Falk, C.; Wenny, M. B.; Mollo, A.; Zeller, M.; Friedler, S. A.; Schrier, J.; Norquist, A. J. Machine-Learning-Assisted Materials Discovery Using Failed Experiments. *Nature* **2016**, *533* (7601), 73–76.
- (32) Ramprasad, R.; Batra, R.; Pilania, G.; Mannodi-Kanakkithodi, A.; Kim, C. Machine Learning in Materials Informatics: Recent Applications and Prospects. *Npj Comput. Mater.* **2017**, *3* (1), 54.
- (33) Isayev, O.; Oses, C.; Toher, C.; Gossett, E.; Curtarolo, S.; Tropsha, A. Universal Fragment Descriptors for Predicting Properties of Inorganic Crystals. *Nat. Commun.* **2017**, *8* (1), 15679.
- (34) Chen, C.; Ye, W.; Zuo, Y.; Zheng, C.; Ong, S. P. Graph Networks as a Universal Machine Learning Framework for Molecules and Crystals. *Chem. Mater.* **2019**, *31* (9), 3564–3572.
- (35) Rupp, M.; Tkatchenko, A.; Müller, K.-R.; Von Lilienfeld, O. A. Fast and Accurate Modeling of Molecular Atomization Energies with Machine Learning. *Phys. Rev. Lett.* **2012**, *108* (5), 058301.
- (36) Hansen, K.; Biegler, F.; Ramakrishnan, R.; Pronobis, W.; Von Lilienfeld, O. A.; Müller, K.-R.; Tkatchenko, A. Machine Learning Predictions of Molecular Properties: Accurate Many-Body Potentials and Nonlocality in Chemical Space. *J. Phys. Chem. Lett.* **2015**, *6* (12), 2326–2331.
- (37) Kearnes, S.; McCloskey, K.; Berndl, M.; Pande, V.; Riley, P. Molecular Graph Convolutions: Moving beyond Fingerprints. *J. Comput. Aided Mol. Des.* **2016**, *30* (8), 595–608.
- (38) Rogers, D.; Hahn, M. Extended-Connectivity Fingerprints. *J. Chem. Inf. Model.* **2010**, *50* (5), 742–754.
- (39) Schütt, K. T.; Arbabzadah, F.; Chmiela, S.; Müller, K. R.; Tkatchenko, A. Quantum-Chemical Insights from Deep Tensor Neural Networks. *Nat. Commun.* **2017**, *8* (1), 13890.
- (40) Schütt, K. T.; Sauceda, H. E.; Kindermans, P.-J.; Tkatchenko, A.; Müller, K.-R. SchNet - A Deep Learning Architecture for Molecules and Materials. *J. Chem. Phys.* **2018**, *148* (24), 241722.
- (41) Goodall, R. E. A.; Lee, A. A. Predicting Materials Properties without Crystal Structure: Deep Representation Learning from Stoichiometry. *Nat. Commun.* **2020**, *11* (1), 6280.
- (42) Park, C. W.; Wolverton, C. Developing an Improved Crystal Graph Convolutional Neural Network Framework for Accelerated Materials Discovery. *Phys. Rev. Mater.* **2020**, *4* (6), 063801.
- (43) Choudhary, K.; DeCost, B. Atomistic Line Graph Neural Network for Improved Materials Property Predictions. *Npj Comput. Mater.* **2021**, *7* (1), 185.
- (44) Chen, C.; Ong, S. P. A Universal Graph Deep Learning Interatomic Potential for the Periodic Table. *Nat. Comput. Sci.* **2022**, *2* (11), 718–728.
- (45) Yamada, H.; Liu, C.; Wu, S.; Koyama, Y.; Ju, S.; Shiomi, J.; Morikawa, J.; Yoshida, R. Predicting Materials Properties with Little Data Using Shotgun Transfer Learning. *ACS Cent. Sci.* **2019**, *5* (10), 1717–1730.
- (46) Magar, R.; Wang, Y.; Barati Farimani, A. Crystal Twins: Self-Supervised Learning for Crystalline Material Property Prediction. *Npj Comput. Mater.* **2022**, *8* (1), 231.
- (47) Dan, Y.; Zhao, Y.; Li, X.; Li, S.; Hu, M.; Hu, J. Generative Adversarial Networks (GAN) Based Efficient Sampling of Chemical Composition Space for Inverse Design of Inorganic Materials. *Npj Comput. Mater.* **2020**, *6* (1), 84.
- (48) Karniadakis, G. E.; Kevrekidis, I. G.; Lu, L.; Perdikaris, P.; Wang, S.; Yang, L. Physics-Informed Machine Learning. *Nat. Rev. Phys.* **2021**, *3* (6), 422–440.
- (49) Agrawal, A.; Choudhary, A. Perspective: Materials Informatics and Big Data: Realization of the “Fourth Paradigm” of Science in Materials Science. *APL Mater.* **2016**, *4* (5), 053208.
- (50) Pyzer-Knapp, E.; Suh, C.; Gómez-Bombarelli, R.; Aguilera-Iparraguirre, J.; Aspuru-Guzik, A. What Is High-Throughput Virtual Screening? A Perspective from Organic Materials Discovery. *Annu. Rev. Mater. Res.* **2015**, *45*, 195–216.
- (51) Merchant, A.; Batzner, S.; Schoenholz, S. S.; Aykol, M.; Cheon, G.; Cubuk, E. D. Scaling Deep Learning for Materials Discovery. *Nature* **2023**, *624* (7990), 80–85.
- (52) Sorkun, M. C.; Astruc, S.; Koelman, J. M. V. A.; Er, S. An Artificial Intelligence-Aided Virtual Screening Recipe for Two-Dimensional Materials Discovery. *Npj Comput. Mater.* **2020**, *6* (1), 106.
- (53) Wang, S.; Liu, C.; Hao, W.; Zhuang, Y.; Zhu, X.; Wang, L.; Niu, X.; Liu, S.; Chen, B.; Zhao, Q. From Formability to Bandgap: Machine Learning Accelerates the Discovery and Application of Perovskite Materials. *ACS Nano* **2025**, *19* (32), 29049–29072.
- (54) Wang, X.; Sheng, Y.; Ning, J.; Xi, J.; Xi, L.; Qiu, D.; Yang, J.; Ke, X. A Critical Review of Machine Learning Techniques on Thermo-electric Materials. *J. Phys. Chem. Lett.* **2023**, *14* (7), 1808–1822.
- (55) Yang, H.; Hu, C.; Zhou, Y.; Liu, X.; Shi, Y.; Li, J.; Li, G.; Chen, Z.; Chen, S.; Zeni, C.; Horton, M.; Pinsler, R.; Fowler, A.; Zügner, D.; Xie, T.; Smith, J.; Sun, L.; Wang, Q.; Kong, L.; Liu, C.; Hao, H.; Lu, Z. MatterSim: A Deep Learning Atomistic Model Across Elements, Temperatures and Pressures. *arXiv* **2024**, DOI: 10.48550/arXiv.2405.04967, May 10, (accessed 2026-04-01).
- (56) Batatia, I.; Benner, P.; Chiang, Y.; Elena, A. M.; Kovács, D. P.; Riebesell, J.; Advincola, X. R.; Asta, M.; Avaylon, M.; Baldwin, W. J.; Berger, F.; Bernstein, N.; Bhowmik, A.; Bigi, F.; Blau, S. M.; Cărare, V.; Ceriotti, M.; Chong, S.; Darby, J. P.; De, S.; Pia, F. D.; Deringer, V. L.; Elijošius, R.; El-Machachi, Z.; Falcioni, F.; Fako, E.; Ferrari, A. C.; Gardner, J. L. A.; Gawkowski, M. J.; Genreith-Schriever, A.; George, J.; Goodall, R. E. A.; Grandel, J.; Grey, C. P.; Grigorev, P.; Han, S.; Handley, W.; Heenen, H. H.; Hermansson, K.; Holm, C.; Ho, C. H.; Hofmann, S.; Jaafar, J.; Jakob, K. S.; Jung, H.; Kapil, V.; Kaplan, A. D.; Karimitari, N.; Kermode, J. R.; Kourtis, P.; Kroupa, N.; Kullgren, J.; Kuner, M. C.; Kuryla, D.; Liepuoniute, G.; Lin, C.; Margraf, J. T.; Magdău, I.-B.; Michaelides, A.; Moore, J. H.; Naik, A. A.; Niblett, S. P.; Norwood, S. W.; O'Neill, N.; Ortner, C.; Persson, K. A.; Reuter, K.; Rosen, A. S.; Rosset, L. A. M.; Schaaf, L. L.; Schran, C.; Shi, B. X.; Sivonxay, E.; Stenczel, T. K.; Svahn, V.; Sutton, C.; Swinburne, T. D.; Tilly, J.; van der Oord, C.; Vargas, S.; Varga-Umbrich, E.; Vegge, T.; Vondrák, M.; Wang, Y.; Witt, W. C.; Wolf, T.; Zills, F.; Csányi, G. A Foundation Model for Atomistic Materials Chemistry. *arXiv* **2025**, DOI: 10.48550/arXiv.2401.00096, September 4, (accessed 2026-04-01).
- (57) Tang, B.; Lu, Y.; Zhou, J.; Chouhan, T.; Wang, H.; Golani, P.; Xu, M.; Xu, Q.; Guan, C.; Liu, Z. Machine Learning-Guided Synthesis of Advanced Inorganic Materials. *Mater. Today* **2020**, *41*, 72–80.
- (58) Segler, M. H. S.; Preuss, M.; Waller, M. P. Planning Chemical Syntheses with Deep Neural Networks and Symbolic AI. *Nature* **2018**, *555* (7698), 604–610.
- (59) Shields, B. J.; Stevens, J.; Li, J.; Parasram, M.; Damani, F.; Alvarado, J. I. M.; Janey, J. M.; Adams, R. P.; Doyle, A. G. Bayesian Reaction Optimization as a Tool for Chemical Synthesis. *Nature* **2021**, *590* (7844), 89–96.
- (60) Shahriari, B.; Swersky, K.; Wang, Z.; Adams, R. P.; de Freitas, N. Taking the Human Out of the Loop: A Review of Bayesian Optimization. *Proc. IEEE* **2016**, *104* (1), 148–175.
- (61) Chang, J.; Nikolaev, P.; Carpena-Núñez, J.; Rao, R.; Decker, K.; Islam, A. E.; Kim, J.; Pitt, M. A.; Myung, J. I.; Maruyama, B. Efficient Closed-Loop Maximization of Carbon Nanotube Growth Rate Using Bayesian Optimization. *Sci. Rep.* **2020**, *10* (1), 9040.
- (62) Attia, P. M.; Grover, A.; Jin, N.; Severson, K. A.; Markov, T. M.; Liao, Y.-H.; Chen, M. H.; Cheong, B.; Perkins, N.; Yang, Z.; Herring, P. K.; Aykol, M.; Harris, S. J.; Braatz, R. D.; Ermon, S.; Chueh, W. C. Closed-Loop Optimization of Fast-Charging Protocols for Batteries with Machine Learning. *Nature* **2020**, *578* (7795), 397–402.
- (63) Paulson, N. H.; Yanguas-Gil, A.; Abuomar, O. Y.; Elam, J. W. Intelligent Agents for the Optimization of Atomic Layer Deposition. *ACS Appl. Mater. Interfaces* **2021**, *13* (14), 17022–17033.
- (64) Huang, X.; Zhao, C. Y.; Wang, H.; Ju, S. AI-Assisted Inverse Design of Sequence-Ordered High Intrinsic Thermal Conductivity Polymers. *Mater. Today Phys.* **2024**, *44*, 101438.

- (65) Batra, R.; Loeffler, T. D.; Chan, H.; Srinivasan, S.; Cui, H.; Korendovych, I. V.; Nanda, V.; Palmer, L. C.; Solomon, L. A.; Fry, H. C.; Sankaranarayanan, S. K. R. S. Machine Learning Overcomes Human Bias in the Discovery of Self-Assembling Peptides. *Nat. Chem.* **2022**, *14* (12), 1427–1435.
- (66) Jiang, Y.; Salley, D.; Sharma, A.; Keenan, G.; Mullin, M.; Cronin, L. An Artificial Intelligence Enabled Chemical Synthesis Robot for Exploration and Optimization of Nanomaterials. *Sci. Adv.* **2022**, *8* (40), No. eabo2626.
- (67) Wu, T.; Zhou, M.; Zou, J.; Chen, Q.; Qian, F.; Kurths, J.; Liu, R.; Tang, Y. AI-Guided Few-Shot Inverse Design of HDP-Mimicking Polymers against Drug-Resistant Bacteria. *Nat. Commun.* **2024**, *15* (1), 6288.
- (68) Zhu, Y.; Fang, J.; Ahmed, S. A. H.; Zhang, T.; Zeng, S.; Liao, J.-Y.; Ma, Z.; Qian, L. A Modular Artificial Intelligence Framework to Facilitate Fluorophore Design. *Nat. Commun.* **2025**, *16* (1), 3598.
- (69) Xie, E.; Wang, X.; Siepmann, J. I.; Chen, H.; Snurr, R. Q. Generative AI for Design of Nanoporous Materials: Review and Future Prospects. *Digit. Discovery* **2025**, *4* (9), 2336–2363.
- (70) Szymanski, N. J.; Nevatia, P.; Bartel, C. J.; Zeng, Y.; Ceder, G. Autonomous and Dynamic Precursor Selection for Solid-State Materials Synthesis. *Nat. Commun.* **2023**, *14* (1), 6956.
- (71) Liao, H.; Hu, S.; Yang, H.; Wang, L.; Tanaka, S.; Takigawa, I.; Li, W.; Fan, H.; Gong, J. P. Data-Driven De Novo Design of Super-Adhesive Hydrogels. *Nature* **2025**, *644*, 89.
- (72) Sohail, Y.; Zhang, C.; Xue, D.; Zhang, J.; Zhang, D.; Gao, S.; Yang, Y.; Fan, X.; Zhang, H.; Liu, G.; Sun, J.; Ma, E. Machine-Learning Design of Ductile FeNiCoAlTa Alloys with High Strength. *Nature* **2025**, *643*, 119.
- (73) Noack, M. M.; Zwart, P. H.; Ushizima, D. M.; Fukuto, M.; Yager, K. G.; Elbert, K. C.; Murray, C. B.; Stein, A.; Doerk, G. S.; Tsai, E. H. R.; Li, R.; Freychet, G.; Zhernenkov, M.; Holman, H.-Y. N.; Lee, S.; Chen, L.; Rotenberg, E.; Weber, T.; Goc, Y. L.; Boehm, M.; Steffens, P.; Mutti, P.; Sethian, J. A. Gaussian Processes for Autonomous Data Acquisition at Large-Scale Synchrotron and Neutron Facilities. *Nat. Rev. Phys.* **2021**, *3* (10), 685–697.
- (74) Harris, S. B.; Vasudevan, R.; Liu, Y. Active Oversight and Quality Control in Standard Bayesian Optimization for Autonomous Experiments. *Npj Comput. Mater.* **2025**, *11* (1), 23.
- (75) Slattery, A.; Wen, Z.; Tenblad, P.; Sanjosé-Orduna, J.; Pintossi, D.; Den Hartog, T.; Noël, T. Automated Self-Optimization, Intensification, and Scale-up of Photocatalysis in Flow. *Science* **2024**, *383* (6681), No. eadj1817.
- (76) Krishnadasan, S.; Brown, R. J. C.; deMello, A. J.; deMello, J. C. Intelligent Routes to the Controlled Synthesis of Nanoparticles. *Lab. Chip* **2007**, *7* (11), 1434–1441.
- (77) Bédard, A.-C.; Adamo, A.; Aroh, K. C.; Russell, M. G.; Bedermann, A. A.; Torosian, J.; Yue, B.; Jensen, K. F.; Jamison, T. F. Reconfigurable System for Automated Optimization of Diverse Chemical Reactions. *Science* **2018**, *361* (6408), 1220–1225.
- (78) Li, Y.; Wang, S.; Lv, Z.; Wang, Z.; Zhao, Y.; Xie, Y.; Xu, Y.; Qian, L.; Yang, Y.; Zhao, Z.; Zhang, J. Transforming the Synthesis of Carbon Nanotubes with Machine Learning Models and Automation. *Matter* **2025**, *8* (1), 101913.
- (79) Zhao, Y.; Liao, J.; Bu, S.; Hu, Z.; Hu, J.; Lu, Q.; Shang, M.; Guo, B.; Chen, G.; Zhao, Q.; Jia, K.; Wang, G.; Errington, E.; Xie, Q.; Zhang, Y.; Guo, M.; Mao, B.; Lin, L.; Liu, Z. Automated Processing and Transfer of Two-Dimensional Materials with Robotics. *Nat. Chem. Eng.* **2025**, *2* (5), 296–308.
- (80) Du, X.; Lüer, L.; Heumueller, T.; Wagner, J.; Berger, C.; Osterrieder, T.; Wortmann, J.; Langner, S.; Vongsaysy, U.; Bertrand, M.; Li, N.; Stubhan, T.; Hauch, J.; Brabec, C. J. Elucidating the Full Potential of OPV Materials Utilizing a High-Throughput Robot-Based Platform and Machine Learning. *Joule* **2021**, *5* (2), 495–506.
- (81) Zhang, B.; Merker, L.; Sanin, A.; Stein, H. S. Robotic Cell Assembly to Accelerate Battery Research. *Digit. Discovery* **2022**, *1* (6), 755–762.
- (82) Ha, T.; Lee, D.; Kwon, Y.; Park, M. S.; Lee, S.; Jang, J.; Choi, B.; Jeon, H.; Kim, J.; Choi, H.; Seo, H.-T.; Choi, W.; Hong, W.; Park, Y. J.; Jang, J.; Cho, J.; Kim, B.; Kwon, H.; Kim, G.; Oh, W. S.; Kim, J. W.; Choi, J.; Min, M.; Jeon, A.; Jung, Y.; Kim, E.; Lee, H.; Choi, Y.-S. AI-Driven Robotic Chemist for Autonomous Synthesis of Organic Molecules. *Sci. Adv.* **2023**, *9* (44), No. eadj0461.
- (83) Burger, B.; Maffettone, P. M.; Gusev, V. V.; Aitchison, C. M.; Bai, Y.; Wang, X.; Li, X.; Alston, B. M.; Li, B.; Clowes, R.; Rankin, N.; Harris, B.; Sprick, R. S.; Cooper, A. I. A Mobile Robotic Chemist. *Nature* **2020**, *583* (7815), 237–241.
- (84) MacLeod, B. P.; Parlani, F. G. L.; Morrissey, T. D.; Häse, F.; Roch, L. M.; Dettelbach, K. E.; Moreira, R.; Yunker, L. P. E.; Rooney, M. B.; Deeth, J. R.; Lai, V.; Ng, G. J.; Situ, H.; Zhang, R. H.; Elliott, M. S.; Haley, T. H.; Dvorak, D. J.; Aspuru-Guzik, A.; Hein, J. E.; Berlinguette, C. P. Self-Driving Laboratory for Accelerated Discovery of Thin-Film Materials. *Sci. Adv.* **2020**, *6* (20), No. eaaz8867.
- (85) 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. An Autonomous Laboratory for the Accelerated Synthesis of Novel Materials. *Nature* **2023**, *624* (7990), 86–91.
- (86) Bromig, L.; Leiter, D.; Mardale, A.-V.; von den Eichen, N.; Bieringer, E.; Weuster-Botz, D. The SiLA 2 Manager for Rapid Device Integration and Workflow Automation. *SoftwareX* **2022**, *17*, 100991.
- (87) Liu, Y.; Roccapiore, K.; Checa, M.; Valletti, S. M.; Yang, J.; Jesse, S.; Vasudevan, R. K. AECrosopy: A Software-Hardware Framework Empowering Microscopy Toward Automated and Autonomous Experimentation. *Small Methods* **2024**, 2301740.
- (88) Coley, C. W.; Barzilay, R.; Jaakkola, T. S.; Green, W. H.; Jensen, K. F. Prediction of Organic Reaction Outcomes Using Machine Learning. *ACS Cent. Sci.* **2017**, *3* (5), 434–443.
- (89) Chen, S.; Jung, Y. Deep Retrosynthetic Reaction Prediction Using Local Reactivity and Global Attention. *JACS Au* **2021**, *1* (10), 1612–1620.
- (90) Ucak, U. V.; Ashynmamatov, I.; Ko, J.; Lee, J. Retrosynthetic Reaction Pathway Prediction through Neural Machine Translation of Atomic Environments. *Nat. Commun.* **2022**, *13* (1), 1186.
- (91) Liu, T.; Cao, Z.; Huang, Y.; Wan, Y.; Wu, J.; Hsieh, C.-Y.; Hou, T.; Kang, Y. SynCluster: Reaction Type Clustering and Recommendation Framework for Synthesis Planning. *JACS Au* **2023**, *3* (12), 3446–3461.
- (92) Han, Y.; Xu, X.; Hsieh, C.-Y.; Ding, K.; Xu, H.; Xu, R.; Hou, T.; Zhang, Q.; Chen, H. Retrosynthesis Prediction with an Iterative String Editing Model. *Nat. Commun.* **2024**, *15* (1), 6404.
- (93) Zhang, P.; Eun, J.; Elkin, M.; Zhao, Y.; Cantrell, R. L.; Newhouse, T. R. A Neural Network Model Informs the Total Synthesis of Clovane Sesquiterpenoids. *Nat. Synth.* **2023**, *2* (6), 527–534.
- (94) Nartova, A. V.; Mashukov, M. Yu.; Astakhov, R. R.; Kudinov, V. Yu.; Matveev, A. V.; Okunev, A. G. Particle Recognition on Transmission Electron Microscopy Images Using Computer Vision and Deep Learning for Catalytic Applications. *Catalysts* **2022**, *12* (2), 135.
- (95) Xu, W.; LeBeau, J. M. A Deep Convolutional Neural Network to Analyze Position Averaged Convergent Beam Electron Diffraction Patterns. *Ultramicroscopy* **2018**, *188*, 59–69.
- (96) Xue, S.; Luo, Z.; Li, A.; Feng, M.; Li, S.; Zhou, S.; Xu, K.; Wang, H.; Zhang, J.; Ouyang, F.; Wang, S. Machine Learning-Assisted Active Center Exploration in Atomically Thin MoS_x Te_{2-x} Electrocatalysts for Efficient Hydrogen Evolution. *Adv. Mater.* **2025**, *37* (39), 2503474.
- (97) Qiu, Z.; Meng, Y.; Li, J.; Hong, Y.; Li, N.; Han, X.; Liang, Y.; Cheng, W. N.; Ke, G.; Zhang, L.; E, W.; Zhao, X.; Zhang, J. Deep Learning for Sub-Ångström-Resolution Imaging in Uncorrected Scanning Transmission Electron Microscopy. *Natl. Sci. Rev.* **2025**, *12* (8), nwaf235.
- (98) Huang, W.; Jin, Y.; Li, Z.; Yao, L.; Chen, Y.; Luo, Z.; Zhou, S.; Lin, J.; Liu, F.; Gao, Z.; Cheng, J.; Zhang, L.; Ouyang, F.; Zhang, J.; Wang, S. Auto-Resolving the Atomic Structure at van Der Waals Interfaces Using a Generative Model. *Nat. Commun.* **2025**, *16* (1), 2927.

- (99) Stuckner, J.; Harder, B.; Smith, T. M. Microstructure Segmentation with Deep Learning Encoders Pre-Trained on a Large Microscopy Dataset. *Npj Comput. Mater.* **2022**, *8* (1), 200.
- (100) Zhang, J.; Zhai, T.; Arifurrahman, F.; Wang, Y.; Hitt, A.; He, Z.; Ai, Q.; Liu, Y.; Lin, C.-Y.; Zhu, Y.; Tang, M.; Lou, J. Toward Controlled Synthesis of 2D Crystals by CVD: Learning from the Real-Time Crystal Morphology Evolutions. *Nano Lett.* **2024**, *24* (8), 2465–2472.
- (101) Smith, B. R.; Pant, B.; Liu, Y.; Liu, Y.-C.; Yang, J.-C.; Jesse, S.; Khojandi, A.; Kalinin, S. V.; Cao, Y.; Vasudevan, R. K. Physics-Informed Models of Domain Wall Dynamics as a Route for Autonomous Domain Wall Design via Reinforcement Learning. *Digit. Discovery* **2024**, *3* (3), 456–466.
- (102) Bulanadi, R.; Chowdhury, J.; Funakubo, H.; Ziatdinov, M.; Vasudevan, R.; Biswas, A.; Liu, Y. Beyond Optimization: Exploring Novelty Discovery in Autonomous Experiments. *ACS Nanosci. Au* **2026**, *6* (1), 86–94.
- (103) Yang, J.; Yin, R. A.; Jiang, C.; Hu, Y.; Zhu, X.; Hu, X.; Kumar, S.; Holmes, S. K.; Wang, X.; Zhai, X.; Rong, K.; Zhu, Y.; Zhang, T.; Yin, Z.; Cao, Y.; Tang, H.; Franklin, A. D.; Kong, J.; Gong, N. Z.; Ren, Z.; Wang, H. Zero-Shot Autonomous Microscopy for Scalable and Intelligent Characterization of 2D Materials. *ACS Nano* **2025**, *19* (40), 35493–35502.
- (104) Kusne, A. G.; Yu, H.; Wu, C.; Zhang, H.; Hattrick-Simpers, J.; DeCost, B.; Sarker, S.; Oses, C.; Toher, C.; Curtarolo, S.; Davydov, A. V.; Agarwal, R.; Bendersky, L. A.; Li, M.; Mehta, A.; Takeuchi, I. On-the-Fly Closed-Loop Materials Discovery via Bayesian Active Learning. *Nat. Commun.* **2020**, *11* (1), S966.
- (105) Salgado, J. E.; Lerman, S.; Du, Z.; Xu, C.; Abdolrahim, N. Automated Classification of Big X-Ray Diffraction Data Using Deep Learning Models. *Npj Comput. Mater.* **2023**, *9* (1), 214.
- (106) Yu, D.; Griesemer, S.; Liu, T.; Wolverton, C.; Zhu, Y. Automated Phase Mapping of High-Throughput X-Ray Diffraction Data Encoded with Domain-Specific Materials Science Knowledge. *Npj Comput. Mater.* **2025**, *11* (1), 354.
- (107) Wei, Y.; Varanasi, R. S.; Schwarz, T.; Gomell, L.; Zhao, H.; Larson, D. J.; Sun, B.; Liu, G.; Chen, H.; Raabe, D.; Gault, B. Machine-Learning-Enhanced Time-of-Flight Mass Spectrometry Analysis. *Patterns* **2021**, *2* (2), 100192.
- (108) Jia, H.; Wang, C.; Wang, C.; Clancy, P. Machine Learning Approach to Enable Spectral Imaging Analysis for Particularly Complex Nanomaterial Systems. *ACS Nano* **2023**, *17* (1), 453–460.
- (109) Zheng, C.; Mathew, K.; Chen, C.; Chen, Y.; Tang, H.; Dozier, A.; Kas, J. J.; Vila, F. D.; Rehr, J. J.; Piper, L. F. J.; Persson, K. A.; Ong, S. P. Automated Generation and Ensemble-Learned Matching of X-Ray Absorption Spectra. *Npj Comput. Mater.* **2018**, *4* (1), 12.
- (110) Brunton, S. L.; Proctor, J. L.; Kutz, J. N. Discovering Governing Equations from Data by Sparse Identification of Nonlinear Dynamical Systems. *Proc. Natl. Acad. Sci. U. S. A.* **2016**, *113* (15), 3932–3937.
- (111) Udrescu, S.-M.; Tegmark, M. AI Feynman: A Physics-Inspired Method for Symbolic Regression. *Sci. Adv.* **2020**, *6* (16), No. eaay2631.
- (112) Wang, T.; Hu, J.; Ouyang, R.; Wang, Y.; Huang, Y.; Hu, S.; Li, W.-X. Nature of Metal-Support Interaction for Metal Catalysts on Oxide Supports. *Science* **2024**, *386* (6724), 915–920.
- (113) Liu, Z.; Tegmark, M. Machine Learning Hidden Symmetries. *Phys. Rev. Lett.* **2022**, *128* (18), 180201.
- (114) Ouyang, R.; Curtarolo, S.; Ahmetcik, E.; Scheffler, M.; Ghiringhelli, L. M. SISSO: A Compressed-Sensing Method for Identifying the Best Low-Dimensional Descriptor in an Immensity of Offered Candidates. *Phys. Rev. Mater.* **2018**, *2* (8), 083802.
- (115) Esterhuizen, J. A.; Goldsmith, B. R.; Linic, S. Uncovering Electronic and Geometric Descriptors of Chemical Activity for Metal Alloys and Oxides Using Unsupervised Machine Learning. *Chem. Catal.* **2021**, *1* (4), 923–940.
- (116) Li, H.; Liu, Y.; Chen, K.; Margraf, J. T.; Li, Y.; Reuter, K. Subgroup Discovery Points to the Prominent Role of Charge Transfer in Breaking Nitrogen Scaling Relations at Single-Atom Catalysts on VS₂. *ACS Catal.* **2021**, *11* (13), 7906–7914.
- (117) Bulmer, J. S.; Sloan, A. W. N.; Glerum, M.; Carpena-Núñez, J.; Waelder, R.; Humes, J.; Boies, A. M.; Pasquali, M.; Rao, R.; Maruyama, B. Forecasting Carbon Nanotube Diameter in Floating Catalyst Chemical Vapor Deposition. *Carbon* **2023**, *201*, 719–733.
- (118) Maulana Kusdhany, M. I.; Lyth, S. M. New Insights into Hydrogen Uptake on Porous Carbon Materials via Explainable Machine Learning. *Carbon* **2021**, *179*, 190–201.
- (119) Feng, J.; Lansford, J. L.; Katsoulakis, M. A.; Vlachos, D. G. Explainable and Trustworthy Artificial Intelligence for Correctable Modeling in Chemical Sciences. *Sci. Adv.* **2020**, *6* (42), No. eabc3204.
- (120) Angello, N. H.; Friday, D. M.; Hwang, C.; Yi, S.; Cheng, A. H.; Torres-Flores, T. C.; Jira, E. R.; Wang, W.; Aspuru-Guzik, A.; Burke, M. D.; Schroeder, C. M.; Diao, Y.; Jackson, N. E. Closed-Loop Transfer Enables Artificial Intelligence to Yield Chemical Knowledge. *Nature* **2024**, *633* (8029), 351–358.
- (121) Tshitoyan, V.; Dagdelen, J.; Weston, L.; Dunn, A.; Rong, Z.; Kononova, O.; Persson, K. A.; Ceder, G.; Jain, A. Unsupervised Word Embeddings Capture Latent Knowledge from Materials Science Literature. *Nature* **2019**, *571* (7763), 95–98.
- (122) Zhou, Q.; Tang, P.; Liu, S.; Pan, J.; Yan, Q.; Zhang, S.-C. Learning Atoms for Materials Discovery. *Proc. Natl. Acad. Sci. U. S. A.* **2018**, *115* (28). DOI: 10.1073/pnas.1801181115.
- (123) Jaeger, S.; Fulle, S.; Turk, S. Mol2vec: Unsupervised Machine Learning Approach with Chemical Intuition. *J. Chem. Inf. Model.* **2018**, *58* (1), 27–35.
- (124) Wei, L.; Li, Q.; Song, Y.; Stefanov, S.; Dong, R.; Fu, N.; Siriwardane, E. M. D.; Chen, F.; Hu, J. Crystal Composition Transformer: Self-Learning Neural Language Model for Generative and Tinkering Design of Materials. *Adv. Sci.* **2024**, *11* (36), 2304305.
- (125) Gupta, T.; Zaki, M.; Krishnan, N. M. A. Mausam. MatSciBERT: A Materials Domain Language Model for Text Mining and Information Extraction. *Npj Comput. Mater.* **2022**, *8* (1), 102.
- (126) Shetty, P.; Rajan, A. C.; Kuenneth, C.; Gupta, S.; Panchumarti, L. P.; Holm, L.; Zhang, C.; Ramprasad, R. A General-Purpose Material Property Data Extraction Pipeline from Large Polymer Corpora Using Natural Language Processing. *Npj Comput. Mater.* **2023**, *9* (1), 52.
- (127) Jia, Y.; Xian, Y.; Xu, Y.; Dang, P.; Ding, X.; Sun, J.; Zhou, Y.; Xue, D. Universal Semantic Embeddings of Chemical Elements for Enhanced Materials Inference and Discovery. *arXiv* **2025**, DOI: 10.48550/arXiv.2502.14912, February 19, (accessed 2026-04-01).
- (128) He, T.; Sun, W.; Huo, H.; Kononova, O.; Rong, Z.; Tshitoyan, V.; Botari, T.; Ceder, G. Similarity of Precursors in Solid-State Synthesis as Text-Mined from Scientific Literature. *Chem. Mater.* **2020**, *32* (18), 7861–7873.
- (129) He, T.; Huo, H.; Bartel, C. J.; Wang, Z.; Cruse, K.; Ceder, G. Precursor Recommendation for Inorganic Synthesis by Machine Learning Materials Similarity from Scientific Literature. *Sci. Adv.* **2023**, *9* (23), No. eadg8180.
- (130) Dagdelen, J.; Dunn, A.; Lee, S.; Walker, N.; Rosen, A. S.; Ceder, G.; Persson, K. A.; Jain, A. Structured Information Extraction from Scientific Text with Large Language Models. *Nat. Commun.* **2024**, *15* (1), 1418.
- (131) Walker, N.; Lee, S.; Dagdelen, J.; Cruse, K.; Gleason, S.; Dunn, A.; Ceder, G.; Alivisatos, A. P.; Persson, K. A.; Jain, A. Extracting Structured Seed-Mediated Gold Nanorod Growth Procedures from Scientific Text with LLMs. *Digit. Discovery* **2023**, *2* (6), 1768–1782.
- (132) Wang, W.; Jiang, X.; Tian, S.; Liu, P.; Dang, D.; Su, Y.; Lookman, T.; Xie, J. Automated Pipeline for Superalloy Data by Text Mining. *Npj Comput. Mater.* **2022**, *8* (1), 9.
- (133) Court, C. J.; Cole, J. M. Magnetic and Superconducting Phase Diagrams and Transition Temperatures Predicted Using Text Mining and Machine Learning. *Npj Comput. Mater.* **2020**, *6* (1), 18.
- (134) Glasby, L. T.; Gubsch, K.; Bence, R.; Oktavian, R.; Isoko, K.; Moosavi, S. M.; Cordiner, J. L.; Cole, J. C.; Moghadam, P. Z.

- DigiMOF: A Database of Metal-Organic Framework Synthesis Information Generated via Text Mining. *Chem. Mater.* **2023**, *35* (11), 4510–4524.
- (135) Choi, J.; Lee, B. Quantitative Topic Analysis of Materials Science Literature Using Natural Language Processing. *ACS Appl. Mater. Interfaces* **2024**, *16* (2), 1957–1968.
- (136) Mehr, S. H. M.; Craven, M.; Leonov, A. I.; Keenan, G.; Cronin, L. A Universal System for Digitization and Automatic Execution of the Chemical Synthesis Literature. *Science* **2020**, *370* (6512), 101–108.
- (137) Pagel, S.; Jirasek, M.; Cronin, L. Validation of the Scientific Literature via Chempuration Augmented by Large Language Models. *arXiv* **2024**, DOI: 10.48550/arXiv.2410.06384, October 8, (accessed 2026-04-01).
- (138) Pei, Z.; Yin, J.; Liaw, P. K.; Raabe, D. Toward the Design of Ultrahigh-Entropy Alloys via Mining Six Million Texts. *Nat. Commun.* **2023**, *14* (1), 54.
- (139) Zaki, M.; Jayadeva; Mausam; Krishnan, N. M. A. MaScQA: Investigating Materials Science Knowledge of Large Language Models. *Digit. Discovery* **2024**, *3* (2), 313–327.
- (140) Jablonka, K. M.; Schwaller, P.; Ortega-Guerrero, A.; Smit, B. Leveraging Large Language Models for Predictive Chemistry. *Nat. Mach. Intell.* **2024**, *6* (2), 161–169.
- (141) Sayeed, H. M.; Smallwood, W.; Baird, S. G.; Sparks, T. D. NLP Meets Materials Science: Quantifying the Presentation of Materials Data in Literature. *Matter* **2024**, *7* (3), 723–727.
- (142) Alampara, N.; Schilling-Wilhelmi, M.; Ríos-García, M.; Mandal, I.; Khetarpal, P.; Grover, H. S.; Krishnan, N. M. A.; Jablonka, K. M. Probing the Limitations of Multimodal Language Models for Chemistry and Materials Research. *Nat. Comput. Sci.* **2025**, *5*, 952.
- (143) Ruan, Y.; Lu, C.; Xu, N.; He, Y.; Chen, Y.; Zhang, J.; Xuan, J.; Pan, J.; Fang, Q.; Gao, H.; Shen, X.; Ye, N.; Zhang, Q.; Mo, Y. An Automatic End-to-End Chemical Synthesis Development Platform Powered by Large Language Models. *Nat. Commun.* **2024**, *15* (1), 10160.
- (144) Zhu, Q.; Zhang, F.; Huang, Y.; Xiao, H.; Zhao, L.; Zhang, X.; Song, T.; Tang, X.; Li, X.; He, G.; Chong, B.; Zhou, J.; Zhang, Y.; Zhang, B.; Cao, J.; Luo, M.; Wang, S.; Ye, G.; Zhang, W.; Chen, X.; Cong, S.; Zhou, D.; Li, H.; Li, J.; Zou, G.; Shang, W.; Jiang, J.; Luo, Y. An All-Round AI-Chemist with a Scientific Mind. *Natl. Sci. Rev.* **2022**, *9* (10), nwac190.
- (145) Zhu, Q.; Huang, Y.; Zhou, D.; Zhao, L.; Guo, L.; Yang, R.; Sun, Z.; Luo, M.; Zhang, F.; Xiao, H.; Tang, X.; Zhang, X.; Song, T.; Li, X.; Chong, B.; Zhou, J.; Zhang, Y.; Zhang, B.; Cao, J.; Zhang, G.; Wang, S.; Ye, G.; Zhang, W.; Zhao, H.; Cong, S.; Li, H.; Ling, L.-L.; Zhang, Z.; Shang, W.; Jiang, J.; Luo, Y. Automated Synthesis of Oxygen-Producing Catalysts from Martian Meteorites by a Robotic AI Chemist. *Nat. Synth.* **2024**, *3* (3), 319–328.
- (146) Darvish, K.; Skreta, M.; Zhao, Y.; Yoshikawa, N.; Som, S.; Bogdanovic, M.; Cao, Y.; Hao, H.; Xu, H.; Aspuru-Guzik, A.; Garg, A.; Shkurti, F. ORGANA: A Robotic Assistant for Automated Chemistry Experimentation and Characterization. *Matter* **2025**, *8* (2), 101897.
- (147) Ren, Z.; Zhang, Z.; Tian, Y.; Li, J. CRESt - Copilot for Real-World Experimental Scientist. *ChemRxiv* **2023**, DOI: 10.26434/chemrxiv-2023-tnz1x, July 11, (accessed 2026-04-01).
- (148) Liu, Y.; Checa, M.; Vasudevan, R. K. Synergizing Human Expertise and AI Efficiency with Language Model for Microscopy Operation and Automated Experiment Design*. *Mach. Learn. Sci. Technol.* **2024**, *5* (2), 02LT01.
- (149) Mathur, S.; Der Vleuten, N. V.; Yager, K. G.; Tsai, E. H. R. VISION: A Modular AI Assistant for Natural Human-Instrument Interaction at Scientific User Facilities. *Mach. Learn. Sci. Technol.* **2025**, *6* (2), 025051.
- (150) Epps, R. W.; Bowen, M. S.; Volk, A. A.; Abdel-Latif, K.; Han, S.; Reyes, K. G.; Amassian, A.; Abolhasani, M. Artificial Chemist: An Autonomous Quantum Dot Synthesis Bot. *Adv. Mater.* **2020**, *32* (30), 2001626.
- (151) Ni, Z.; Li, Y.; Hu, K.; Han, K.; Xu, M.; Chen, X.; Liu, F.; Ye, Y.; Bai, S. MatPilot: An LLM-Enabled AI Materials Scientist under the Framework of Human-Machine Collaboration. *arXiv* **2024**, DOI: 10.48550/arXiv.2411.08063 (accessed 2026-04-01).
- (152) Biswas, A.; Liu, Y.; Creange, N.; Liu, Y.-C.; Jesse, S.; Yang, J.-C.; Kalinin, S. V.; Ziatdinov, M. A.; Vasudevan, R. K. A Dynamic Bayesian Optimized Active Recommender System for Curiosity-Driven Partially Human-in-the-Loop Automated Experiments. *Npj Comput. Mater.* **2024**, *10* (1), 29.
- (153) Pratiush, U.; Funakubo, H.; Vasudevan, R.; Kalinin, S. V.; Liu, Y. Scientific Exploration with Expert Knowledge (SEEK) in Autonomous Scanning Probe Microscopy with Active Learning. *Digit. Discovery* **2025**, *4* (1), 252–263.
- (154) Zhou, J.; Zhang, B.; Li, G.; Chen, X.; Li, H.; Xu, X.; Chen, S.; He, W.; Xu, C.; Liu, L.; Gao, X. An AI Agent for Fully Automated Multi-Omic Analyses. *Adv. Sci.* **2024**, *11* (44), 2407094.
- (155) First, E.; Rabe, M. N.; Ringer, T.; Brun, Y. Baldur: Whole-Proof Generation and Repair with Large Language Models. In *Proceedings of the 31st ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering*; ACM: San Francisco CA USA, 2023; pp 1229–1241. DOI: 10.1145/3611643.3616243.
- (156) O'Donoghue, O.; Shtedritski, A.; Ginger, J.; Abboud, R.; Ghareeb, A.; Rodrigues, S. BioPlanner: Automatic Evaluation of LLMs on Protocol Planning in Biology. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*; Association for Computational Linguistics: Singapore, 2023; pp 2676–2694. DOI: 10.18653/v1/2023.emnlp-main.162.
- (157) Kim, S.; Jung, Y.; Schrier, J. Large Language Models for Inorganic Synthesis Predictions. *J. Am. Chem. Soc.* **2024**, *146* (29), 19654–19659.
- (158) Oche, A. J.; Biswas, A. Role of Large Language Models and Retrieval-Augmented Generation for Accelerating Crystalline Material Discovery: A Systematic Review. *arXiv* **2025**, DOI: 10.48550/arXiv.2508.06691, August 8, (accessed 2026-04-01).
- (159) Jia, S.; Zhang, C.; Fung, V. LLMatDesign: Autonomous Materials Discovery with Large Language Models. *arXiv* **2024**, DOI: 10.48550/arXiv.2406.13163, June 19, (accessed 2026-04-01).
- (160) Chen, C.; Maqsood, A.; Zhang, Z.; Wang, X.; Duan, L.; Wang, H.; Chen, T.; Liu, S.; Li, Q.; Luo, J.; Jacobsson, T. J. The Use of ChatGPT to Generate Experimentally Testable Hypotheses for Improving the Surface Passivation of Perovskite Solar Cells. *Cell Rep. Phys. Sci.* **2024**, *5* (7), 102058.
- (161) Liu, S.; Wen, T.; Ye, B.; Li, Z.; Liu, H.; Ren, Y.; Srolovitz, D. J. Large Language Models for Material Property Predictions: Elastic Constant Tensor Prediction and Materials Design. *Digit. Discovery* **2025**, *4* (6), 1625–1638.
- (162) Zhao, Z.; Huang, Z.; Li, J.; Lin, S.; Zhou, J.; Cao, F.; Zhou, K.; Ge, R.; Long, T.; Zhu, Y.; Liu, Y.; Zheng, J.; Wei, J.; Zhu, R.; Zou, P.; Li, W.; Cheng, Z.; Ding, T.; Wang, Y.; Yan, Y.; Wei, T.; Ming, H.; Mao, W.; Sun, C.; Liu, Y.; Wang, Z.; Zhang, Z.; Yang, T.; Ma, H.; Gao, Z.; Pei, J. SUPERChem: A Multimodal Reasoning Benchmark in Chemistry. *arXiv* **2025**, DOI: 10.48550/arXiv.2512.01274 (accessed 2026-04-01).
- (163) 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.; Yang, H.; Hao, H.; Li, J.; Yang, C.; Li, W.; Tomioka, R.; Xie, T. A Generative Model for Inorganic Materials Design. *Nature* **2025**, *639* (8055), 624–632.
- (164) Gupta, A.; Savarese, S.; Ganguli, S.; Fei-Fei, L. Embodied Intelligence via Learning and Evolution. *Nat. Commun.* **2021**, *12* (1), 5721.
- (165) Waldrop, M. M. The Chips Are down for Moore's Law. *Nature* **2016**, *530* (7589), 144–147.
- (166) Theis, T. N.; Wong, H.-S. P. The End of Moore's Law: A New Beginning for Information Technology. *Comput. Sci. Eng.* **2017**, *19* (2), 41–50.

- (167) Novoselov, K. S.; Fal'ko, V. I.; Colombo, L.; Gellert, P. R.; Schwab, M. G.; Kim, K. A Roadmap for Graphene. *Nature* **2012**, *490* (7419), 192–200.
- (168) Novoselov, K. S.; Mishchenko, A.; Carvalho, A.; Castro Neto, A. H. 2D Materials and van Der Waals Heterostructures. *Science* **2016**, *353* (6298), aac9439.
- (169) Wang, J.; Cheng, C.; Zheng, X.; Idrobo, J. C.; Lu, A.-Y.; Park, J.-H.; Shin, B. G.; Jung, S. J.; Zhang, T.; Wang, H.; Gao, G.; Shin, B.; Jin, X.; Ju, L.; Han, Y.; Li, L.-J.; Karnik, R.; Kong, J. Cascaded Compression of Size Distribution of Nanopores in Monolayer Graphene. *Nature* **2023**, *623* (7989), 956–963.
- (170) Hills, G.; Lau, C.; Wright, A.; Fuller, S.; Bishop, M. D.; Srimani, T.; Kanhaiya, P.; Ho, R.; Amer, A.; Stein, Y.; Murphy, D.; Arvind; Chandrakasan, A.; Shulaker, M. M. Modern Microprocessor Built from Complementary Carbon Nanotube Transistors. *Nature* **2019**, *572* (7771), 595–602.
- (171) Liu, L.; Han, J.; Xu, L.; Zhou, J.; Zhao, C.; Ding, S.; Shi, H.; Xiao, M.; Ding, L.; Ma, Z.; Jin, C.; Zhang, Z.; Peng, L.-M. Aligned, High-Density Semiconducting Carbon Nanotube Arrays for High-Performance Electronics. *Science* **2020**, *368* (6493), 850–856.
- (172) Xie, Y.; Li, Y.; Peng, Z.; Wang, C.; Qiu, Z.; Cai, X.; Song, T.; Si, J.; Zhao, X.; Qian, L.; Zhao, Z.; Zhang, J. Nano-Seeding Catalysts for High-Density Arrays of Horizontally Aligned Carbon Nanotubes with Wafer-Scale Uniformity. *Nat. Commun.* **2025**, *16* (1), 149.
- (173) Qin, B.; Jiang, J.; Wang, L.; Guo, Q.; Zhang, C.; Xu, L.; Ni, X.; Yin, P.; Peng, L.-M.; Wang, E.; Ding, F.; Qiu, C.; Liu, C.; Liu, K. Two-Dimensional Indium Selenide Wafers for Integrated Electronics. *Science* **2025**, *389* (6757), 299–302.
- (174) Akinwande, D.; Huyghebaert, C.; Wang, C.-H.; Serna, M. I.; Goossens, S.; Li, L.-J.; Wong, H.-S. P.; Koppens, F. H. L. Graphene and Two-Dimensional Materials for Silicon Technology. *Nature* **2019**, *573* (7775), 507–518.
- (175) Xia, Q.; Yang, J. J. Memristive Crossbar Arrays for Brain-Inspired Computing. *Nat. Mater.* **2019**, *18* (4), 309–323.
- (176) Si, J.; Zhang, P.; Zhao, C.; Lin, D.; Xu, L.; Xu, H.; Liu, L.; Jiang, J.; Peng, L.-M.; Zhang, Z. A Carbon-Nanotube-Based Tensor Processing Unit. *Nat. Electron.* **2024**, *7* (8), 684–693.
- (177) Bandyopadhyay, S.; Sludds, A.; Krastanov, S.; Hamerly, R.; Harris, N.; Bunandar, D.; Streshinsky, M.; Hochberg, M.; Englund, D. Single-Chip Photonic Deep Neural Network with Forward-Only Training. *Nat. Photonics* **2024**, *18* (12), 1335–1343.
- (178) Jeon, K.; Ryu, J. J.; Im, S.; Seo, H. K.; Eom, T.; Ju, H.; Yang, M. K.; Jeong, D. S.; Kim, G. H. Purely Self-Rectifying Memristor-Based Passive Crossbar Array for Artificial Neural Network Accelerators. *Nat. Commun.* **2024**, *15* (1), 129.
- (179) Zhan, Z.; Gao, Y.; Liao, Y.; Xie, W.; Liu, S.; Wang, X. Perovskite Neuromorphic Engine for Transformer Architectures. *Adv. Sci.* **2025**, *12* (33), No. e04706.
- (180) Rus, D.; Tolley, M. T. Design, Fabrication and Control of Soft Robots. *Nature* **2015**, *521* (7553), 467–475.
- (181) Howard, D.; Eiben, A. E.; Kennedy, D. F.; Mouret, J.-B.; Valencia, P.; Winkler, D. Evolving Embodied Intelligence from Materials to Machines. *Nat. Mach. Intell.* **2019**, *1* (1), 12–19.
- (182) Bao, Z. Skin-Inspired Organic Electronic Materials and Devices. *MRS Bull.* **2016**, *41* (11), 897–904.
- (183) Liu, Y.; Zhu, Y.; Cui, Y. Challenges and Opportunities towards Fast-Charging Battery Materials. *Nat. Energy* **2019**, *4* (7), 540–550.
- (184) Liu, C.; Li, F.; Ma, L.-P.; Cheng, H.-M. Advanced Materials for Energy Storage. *Adv. Mater.* **2010**, *22* (8), No. E28-E62.
- (185) Jena, A. K.; Kulkarni, A.; Miyasaka, T. Halide Perovskite Photovoltaics: Background, Status, and Future Prospects. *Chem. Rev.* **2019**, *119* (5), 3036–3103.
- (186) Stranks, S. D.; Snaith, H. J. Metal-Halide Perovskites for Photovoltaic and Light-Emitting Devices. *Nat. Nanotechnol.* **2015**, *10* (5), 391–402.
- (187) Allen, T.; Busby, J.; Meyer, M.; Petti, D. Materials Challenges for Nuclear Systems. *Mater. Today* **2010**, *13* (12), 14–23.