

REVIEW

Open access

Artificial Intelligence as a Co-Scientist in Materials Science: From Pattern Recognition to Self-Driving Laboratories

Jose Martinez^{1*}, Carmen Lopez¹

Abstract

Artificial intelligence is rapidly moving beyond its early role as a pattern-recognition and predictive-modelling tool in materials science. What began as an acceleration strategy for screening known datasets is now becoming a broader transformation of how materials hypotheses are generated, tested, and refined. The central problem is that this transformation is often described in fragments: predictive models in one literature, generative design in another, physics-informed learning in another, and autonomous laboratories in yet another. A unified conceptual synthesis is needed to explain how these streams collectively move AI from passive assistant to active scientific collaborator. This integrative review traces the evolution of AI in materials science from 2017 to 2026. It frames the field through the idea of the AI co-scientist: an intelligent system that can recognise patterns, propose candidates, incorporate physical constraints, select experiments, and learn from feedback. The review integrates 31 peer-reviewed articles spanning materials informatics, machine learning, generative AI, inverse design, physics-informed modelling, active learning, autonomous experimentation, and self-driving laboratories. It does not present new empirical data, meta-analysis, or bibliometric mapping. The synthesis identifies four major evolutionary stages: pattern recognition, generative design, physics-integrated AI, and autonomous experimentation. These stages are not isolated phases but mutually reinforcing capabilities that increasingly connect computation, synthesis, characterisation, and human judgement. The review concludes that AI is becoming a genuine partner in materials discovery, but this transition depends on trustworthy data infrastructure, interpretable models, robust experimental integration, and new norms for human–AI collaboration. The co-scientist paradigm offers a forward-looking framework for understanding how materials science may be reorganised around closed-loop intelligence.

Keywords Artificial intelligence, Materials science, Self-driving laboratories, Co-scientist, Pattern recognition, Generative AI

*Correspondence:

Jose Martinez
jose.martinez@gmail.com

¹ Department of Intelligent Materials Engineering, University of Salamanca, Salamanca, Spain

Introduction

Materials discovery has traditionally relied on a slow sequence of intuition, synthesis, characterisation, theory, and iteration. This workflow has produced transformative materials, but it is constrained by the size of chemical and structural search spaces, the cost of experiments, and the difficulty of connecting composition, processing, structure, and properties. The Materials Genome Initiative and related computational efforts attempted to accelerate this cycle by linking high-throughput computation, curated databases, and materials design logic [1]. Within this setting, AI emerged as a way to extract structure from expanding experimental and computational datasets rather than relying only on human-guided trial and error.

The first influential wave of AI in materials science treated machine learning mainly as a predictive accelerator. Reviews by Ramprasad, Batra, Pilania, Mannodi-Kanakkithodi, and Kim framed materials informatics as a route to property prediction, screening, and rational selection of candidates from large datasets [2]. Butler, Davies, Cartwright, Isayev, and Walsh similarly positioned machine learning as a general approach for molecular and materials science, capable of learning relationships that are difficult to encode directly through handcrafted physical models [3]. In this early framing, AI enhanced the speed and scope of materials analysis but remained dependent on human scientists to define questions, choose descriptors, and interpret outputs.

The more recent transformation is that AI systems increasingly do more than classify, regress, or rank known candidates. Generative models can propose molecular or crystalline structures, active-learning systems can choose what experiment should be performed next, and autonomous laboratories can execute parts of the discovery cycle without continuous human intervention [4-6]. These developments shift AI from a passive pattern recogniser toward an active participant in the formulation and testing of scientific possibilities. The conceptual challenge is therefore to understand not only whether AI improves prediction accuracy, but how it changes the organisation of scientific work.

This integrative review addresses that challenge by synthesising predictive modelling, generative inverse design, physics-informed learning, and self-driving laboratories into a single evolutionary narrative. It argues that the central movement in the field is from AI as a tool for analysing materials data to AI as a co-scientist that participates in hypothesis generation, experimental planning, and interpretation. The review draws on advances in graph neural networks, learned representations, generative design, physics-informed machine learning, and autonomous experimentation to build this argument [7-10]. Its purpose is conceptual integration rather than systematic enumeration, with emphasis on how previously separate technical streams converge into the co-scientist paradigm.

Scope and Evolution of AI in Materials Science

The scope of this review includes AI methods that intervene at different points in the materials discovery cycle: prediction, design, constraint enforcement, experiment selection, robotic execution, and interpretation. Predictive models such as composition-based networks and graph neural networks learn relationships between materials representations and properties, while generative models search for new candidates with desired functions [4, 7, 11]. Physics-informed methods introduce constraints that make learned outputs more consistent with governing equations or atomistic behaviour [10, 12]. Autonomous experimentation then embeds these computational capabilities into closed-loop workflows that connect algorithms with instruments and synthesis platforms [4, 5].

The field can be understood as an expanding ladder of agency. At the first level, AI recognises patterns in existing data and predicts likely properties; at the second, it proposes new candidates through inverse design; at the third, it reasons within physical constraints; and at the fourth, it helps select, run, and interpret experiments in closed loops. This progression is visible in the movement from materials informatics reviews focused on prediction [2, 13] to self-driving laboratory frameworks that emphasise automated discovery cycles [5, 14]. The co-scientist idea emerges when these capabilities are integrated rather than treated as isolated tools.

Figure 1 illustrates the staged evolution of artificial intelligence from pattern recognition to an integrated co-scientist architecture in materials discovery.




Figure 1. Evolutionary Architecture of the AI Co-Scientist in Materials Science

The evolution also reflects a shift in representation. Early models often relied on engineered descriptors, but later methods increasingly learned from composition, crystal graphs, stoichiometry, text, and experimental feedback [7, 8, 15, 16]. Graph networks made crystal and molecular structures computable as relational systems, while unsupervised language models demonstrated that latent materials knowledge could be extracted from scientific literature itself [8, 15]. These representational advances widened the kinds of knowledge that AI could use, making it possible to connect formal databases with less structured scientific information.

Table 1 maps the key evolutionary stages of AI in materials science, from data analysis to co-scientist. The table organises the review’s conceptual framework by linking each stage to its dominant AI capability, scientific role, and representative developments in the literature [2, 4-6, 10]. It also shows that the co-scientist paradigm is not a single technology but a layered integration of prediction, generation, physical reasoning, and autonomous action. This staged view clarifies why the most important question is no longer simply how accurate AI models are, but how responsibly and productively they can participate in discovery.

Table 1. Evolution of AI in Materials Science: Stages, Capabilities, and Landmark Developments

Evolutionary stage	Dominant AI capability	Main scientific function	Representative developments	Co-scientist implication
Pattern recognition	Supervised learning, descriptor-based models, random forests, neural networks	Predict properties and rank candidates from existing datasets	Materials informatics frameworks, composition-based learning, early neural networks [2, 11, 13]	AI accelerates screening but remains mostly passive
Structure-aware prediction	Graph neural networks and learned representations	Encode crystal, molecular, and stoichiometric structure	Crystal graph convolutional networks, universal graph networks, stoichiometric representation learning [7, 15, 16]	AI begins to learn chemically meaningful representations
Generative inverse design	Variational autoencoders, generative adversarial networks, reinforcement learning	Propose new molecules, compositions, and crystal structures	Matter engineering through generative models and constrained crystal generation [4, 9, 17]	AI becomes a design partner rather than only a predictor

Physics-integrated AI	Physics-informed neural networks and hybrid atomistic models	Constrain predictions and generated candidates with physical laws	Physically informed atomistic modelling and physics-informed machine learning [10, 12, 18]	AI outputs become more scientifically grounded and testable
Self-driving laboratories	Active learning, robotic synthesis, automated characterisation	Close the loop between hypothesis, experiment, and feedback	Thin-film self-driving laboratories, mobile robotic chemists, autonomous chemistry platforms [6, 14, 19]	AI participates in experiment selection and iterative discovery
Co-scientist paradigm	Integrated prediction, generation, reasoning, experimentation, and interpretation	Collaborate with humans across the discovery cycle	Emerging materials intelligence ecosystems and autonomous discovery roadmaps [14, 20]	AI becomes a scientific collaborator under human direction

Pattern Recognition and Predictive Modeling

Pattern recognition was the foundational mode through which AI entered mainstream materials science. The key idea was that material properties could be inferred from patterns in prior computational or experimental data, reducing the need to test every candidate directly. Ramprasad, Batra, Pilania, Mannodi-Kanakithodi, and Kim described this approach as a materials informatics strategy in which algorithms learn from descriptors that encode composition, structure, or processing histories [2]. The strength of this phase was speed, but its limitation was that it mostly answered questions posed by human investigators.

Early predictive models depended heavily on how materials were represented. ElemNet showed that deep learning could infer chemical trends directly from elemental composition, reducing dependence on elaborate handcrafted descriptors while retaining predictive power [11]. Stoichiometry-based representation learning later extended this idea by predicting materials properties even when detailed crystal structures were unavailable [16]. These models were important because many real discovery problems begin with incomplete structural information, especially when screening broad compositional spaces.

Graph-based learning marked a major step beyond descriptor engineering because it represented materials as relational objects rather than flat feature vectors. Crystal graph convolutional neural networks treated atoms and bonds as graph elements, enabling property prediction in a way that was both accurate and more interpretable than many black-box alternatives [7]. Graph networks were then generalised as a framework for molecules and crystals, helping unify molecular and solid-state prediction under a common representation [15]. This development began to make AI more chemically fluent, because the model architecture reflected the connectivity and neighbourhood structure of matter.

Predictive modelling nevertheless remained bounded by the data and hypotheses available to it. Reviews of solid-state machine learning emphasised that models could accelerate screening but also inherited biases from training datasets, descriptor choices, and uneven coverage of materials classes [13]. The broader materials intelligence ecosystem proposed by Batra, Song, and Ramprasad recognised that prediction must be connected to data curation, uncertainty estimation, automation, and decision-making rather than treated as an isolated modelling exercise [20]. In this sense, pattern recognition was necessary but insufficient: it built the computational foundation for AI-enabled discovery, but not yet the full agency associated with a co-scientist.

Generative AI and Inverse Design

Generative AI changed the role of AI from evaluation to proposal. Instead of asking whether a known material has a desired property, inverse design asks what material could satisfy a target function. Sanchez-Lengeling and Aspuru-Guzik framed generative models as tools for matter engineering, using learned latent spaces to generate candidate molecules or materials that may not already exist in a database [4]. This shift is conceptually profound because the algorithm begins to participate in the creative act of proposing scientific possibilities.

In materials science, generative models operate across different representational levels, including molecular graphs, chemical compositions, crystal structures, and continuous latent vectors. Noh, Kim, Stein, Sanchez-Lengeling, Gregoire, Aspuru-Guzik, and Jung demonstrated that solid-state materials could be explored through continuous representations, making inverse design more tractable for discrete and constrained material spaces [17]. Long and colleagues extended this direction through a constrained crystal deep convolutional generative adversarial network for inverse crystal-structure design [9]. These studies show that generative AI can explore candidate spaces that are too large for exhaustive enumeration.

Generative models also expose the difference between novelty and scientific usefulness. A generated candidate is valuable only if it is chemically plausible, synthesizable, stable, and aligned with a target property or function. Reviews of inverse materials design and generative approaches have therefore stressed the need to connect generation with screening, validation, and optimisation rather than treating novelty alone as success [21, 22]. Park and Jung’s question of whether generative AI has solved inverse materials design captures this unresolved tension: generation is powerful, but the path from generated structure to realised material remains difficult [23].

Table 2 summarises the major generative AI architectures and their application to materials design. The table highlights how variational autoencoders, generative adversarial networks, reinforcement-learning strategies, and continuous representations contribute differently to inverse design, while also showing their dependence on constraints, validation, and downstream testing [4, 9, 17, 21-23]. It makes clear that generative AI is best understood as a design partner whose proposals must be filtered through physics, synthesis feasibility, and experimental feedback. This is where generative design begins to connect naturally with physics-informed modelling and autonomous laboratories.

Table 2. Generative AI Architectures for Inverse Materials Design: Models, Inputs, Outputs, and Example Discoveries

Generative architecture	Typical inputs	Design outputs	Strengths for materials discovery	Key limitations
Variational autoencoders	Molecular strings, compositions, structural descriptors, latent representations	New candidates sampled from continuous latent spaces	Enable smooth exploration and optimisation of candidate spaces [4, 17]	May generate invalid or difficult-to-synthesise candidates
Generative adversarial networks	Crystal grids, images, structural encodings, composition–structure data	Candidate crystal structures or material-like representations	Useful for learning distributions of plausible structures [9]	Training instability and physical validity remain concerns
Reinforcement-learning design	Reward functions based on target properties, stability, or constraints	Optimised molecular or materials candidates	Aligns generation with explicit design objectives [4, 22]	Reward design can bias outputs or overfit narrow objectives
Continuous representation models	Encoded solid-state materials spaces	Interpolated or optimised candidate materials	Helps navigate discrete solid-state search spaces [17]	Requires careful decoding and validation

Hybrid generative-screening workflows	Generated candidates plus predictive filters	Ranked candidate sets for further simulation or experiment	Connects creativity with evaluation and prioritisation [21, 22]	Screening quality depends on predictive model reliability
Critical inverse-design frameworks	Literature synthesis, model benchmarking, validation criteria	Roadmaps for realistic deployment	Clarify the gap between generation and real discovery [23]	Reveal unresolved challenges in synthesis and generalisation

Integration with Physics-Based Models

The integration of physics with AI responds to a central weakness of purely data-driven modelling: high statistical accuracy does not guarantee physical realism. Physics-informed neural networks and related approaches embed equations, constraints, symmetries, conservation principles, or atomistic knowledge into learning systems. Raissi, Perdikaris, and Karniadakis introduced physics-informed neural networks as a framework for solving forward and inverse problems involving nonlinear partial differential equations [18]. Although the original formulation was broad, its influence in materials science lies in showing how learning can be constrained by the structure of physical law.

In atomistic materials modelling, physically informed neural networks help bridge the gap between empirical machine learning and first-principles reasoning. Pun, Batra, Ramprasad, and Mishin demonstrated how artificial neural networks can incorporate physical information for atomistic modelling, making learned potentials and predictions more consistent with known material behaviour [12]. This kind of hybridisation is crucial because materials discovery depends not only on prediction but also on whether proposed configurations obey energetics, bonding, stability, and thermodynamic plausibility. Physics-informed learning therefore increases trust by making AI outputs more compatible with scientific reasoning.

Physics-informed machine learning also reframes AI as a reasoning partner rather than a black-box interpolator. Karniadakis, Kevrekidis, Lu, Perdikaris, Wang, and Yang described physics-informed machine learning as a broad paradigm that combines data, models, and governing principles to solve problems where neither data nor theory alone is sufficient [10]. In materials science, this is especially relevant because datasets are often sparse, heterogeneous, and expensive to expand experimentally. By encoding physical constraints, AI can generalise more responsibly across regimes where purely statistical models may fail.

Recent reviews of physics-informed neural networks in materials modelling and design show that this area is becoming a bridge between computational mechanics, materials simulation, and AI-driven discovery [24]. The significance of this bridge is not merely technical; it changes the epistemic status of AI-generated suggestions by linking them to known physical structure. When generative models are constrained by physics and coupled to predictive uncertainty, their outputs become more suitable for experimental prioritisation. This integration is a necessary step toward the AI co-scientist, because a scientific collaborator must not only imagine candidates but reason within the rules of matter.

Self-Driving Laboratories and Autonomous Experimentation

Self-driving laboratories represent the point at which AI moves from computational recommendation to experimental action. In these systems, algorithms select experiments, robotic platforms perform synthesis or testing, automated instruments generate new data, and active-learning models update the next decision cycle. Häse, Roch, and Aspuru-Guzik described this next-generation experimentation model as a route toward laboratories that learn from their own results rather than merely executing fixed human-designed protocols [5]. The conceptual shift is that AI becomes embedded in the material production loop, not confined to offline prediction.

Early demonstrations showed how autonomous systems could compress discovery cycles by linking decision-making algorithms with real experimental hardware. MacLeod and colleagues developed a self-driving laboratory for thin-film materials, showing how automated synthesis, characterisation, and optimisation could be coordinated to accelerate the search for functional materials [6]. Burger and colleagues extended this vision through a mobile robotic chemist capable of navigating a laboratory environment and performing experimental tasks [19]. These platforms made autonomy tangible by showing that AI-guided experimentation could operate in physical laboratory space.

Autonomous experimentation also depends on the logic of active learning, because not every possible experiment should be run. Bayesian optimisation and related strategies allow systems to balance exploration of uncertain regions with exploitation of promising candidates, as illustrated by Shields and colleagues in chemical synthesis optimisation [25]. In materials discovery, this logic is essential because experiments are expensive, noisy, and sometimes irreversible. The self-driving laboratory therefore acts not only as a faster laboratory but as an adaptive scientific system that decides what information is worth acquiring next.

The broader literature on self-driving laboratories shows that autonomy is becoming a unifying architecture for chemistry and materials science. Soldatov, Butova, Guda, Lomachenko, and Lamberti framed self-driving laboratories as platforms for developing new functional materials, while Tom and colleagues presented a comprehensive Chemical Reviews synthesis of the field's methods, challenges, and opportunities [14, 26]. These reviews make clear that the autonomous laboratory is not simply a robot attached to an algorithm; it is a closed-loop discovery ecosystem. In such an ecosystem, prediction, generation, synthesis, characterisation, and learning are connected into a recursive workflow.

Human–AI Collaboration as Co-Scientist

The idea of AI as a co-scientist emerges when predictive, generative, physics-informed, and autonomous capabilities are integrated into a collaborative research workflow. In this model, AI does not replace the scientist but extends the scientist's capacity to search, infer, design, and iterate across vast materials spaces. The materials intelligence ecosystem described by Batra, Song, and Ramprasad provides an important foundation for this view because it links data, models, automation, and decision-making into a coordinated discovery infrastructure [20]. The co-scientist is therefore best understood as a system-level role rather than a single algorithmic function.

Human–AI collaboration begins with complementary strengths. AI can identify correlations in high-dimensional data, generate candidate structures, prioritise experiments, and update models quickly after new evidence appears. Human researchers remain essential for problem framing, mechanistic interpretation, value judgement, safety assessment, and deciding which scientific questions matter. The latent-knowledge extraction demonstrated by Tshitoyan and colleagues illustrates this complementarity because AI can surface patterns from literature, but humans must still interpret their significance within broader scientific contexts [8].

The co-scientist model also requires interaction modes that are more dialogic than traditional software use. A human scientist might ask an AI system to propose hypotheses, compare mechanisms, suggest experiments, or explain why a candidate was prioritised. Explainable machine learning is therefore central, because collaboration depends on the ability to interrogate model reasoning rather than merely accept outputs. Zhong, Gallagher, Liu, Kailkhura, Hiszpanski, and Han emphasised that explainability in materials science is necessary for trust, diagnosis, and the conversion of predictive models into scientific understanding [27].

Table 3 presents the framework for human–AI collaboration and the division of scientific labour. The table positions the AI co-scientist as a collaborator that supports hypothesis generation, candidate design, experimental planning, closed-loop optimisation, and interpretation while leaving strategic judgement and accountability with human researchers [5, 14, 20, 27]. This division of labour avoids two extremes: treating AI as a mere calculator or imagining it as an autonomous replacement for scientific expertise. The most credible future is a negotiated partnership in which AI expands the reachable search space and humans guide meaning, responsibility, and purpose.

Table 3. The Human–AI Co-Scientist Model: Roles, Tasks, and Interaction Modes in Autonomous Materials Discovery

Discovery activity	AI role	Human role	Interaction mode	Co-scientist value
Problem framing	Analyse prior data and literature to identify opportunity spaces	Define scientific goals, constraints, and societal relevance	Human-guided query, AI-supported synthesis	Expands awareness of possible research directions
Hypothesis generation	Propose candidate mechanisms, structures, or compositions	Judge plausibility and refine hypotheses	Iterative dialogue and model interrogation	Converts data patterns into testable ideas
Candidate design	Generate and rank materials with target properties	Set design objectives and evaluate feasibility	Generative proposal with human filtering	Broadens search beyond human intuition
Experimental planning	Select informative experiments using uncertainty and active learning	Approve protocols, manage safety, and interpret trade-offs	Closed-loop recommendation and oversight	Reduces experimental waste and accelerates learning
Robotic execution	Coordinate synthesis, measurement, and feedback capture	Maintain instrumentation and validate protocols	Autonomous execution with human supervision	Connects computation to real materials
Interpretation	Explain model decisions, compare outcomes, and update priors	Build mechanistic explanations and decide next research direction	Explainable AI plus expert reasoning	Supports scientific understanding, not only optimisation
Governance	Track provenance, uncertainty, and model limitations	Ensure ethical, legal, and intellectual accountability	Auditable decision records	Preserves responsibility in autonomous discovery

Figure 2 presents the human–AI closed-loop discovery model through which AI systems and human researchers jointly generate, test, interpret, and refine materials hypotheses.




Figure 2. Human–AI Closed-Loop Discovery Model for Co-Scientific Materials Research

The co-scientist metaphor must nevertheless be used carefully. AI systems can generate plausible suggestions, but they do not possess scientific judgement in the human sense, and their apparent creativity depends on training data, objectives, representations, and constraints. Pilania's discussion of movement from explainable predictions to autonomous design highlights the need to connect model performance with interpretability and design reasoning [28]. A responsible co-scientist framework therefore treats AI as an active collaborator whose outputs remain subject to human scrutiny, physical validation, and experimental evidence.

Challenges and Barriers

The first major barrier is data quality. Materials datasets are often sparse, inconsistent, biased toward successful results, and unevenly distributed across compositions, processing conditions, and property classes. Predictive models trained on such data may appear accurate within benchmark settings but fail when transferred to new material families or experimental environments. Ramprasad and colleagues' materials informatics perspective already made clear that data representation and curation are foundational to reliable learning [2].

The second barrier is interpretability and trust. Deep neural networks, graph models, and generative architectures can make accurate predictions or compelling proposals without clearly communicating why a material was selected. Explainability is particularly important in materials science because scientists seek mechanisms, not only rankings or numerical outputs [27]. Without interpretable reasoning, AI risks becoming a high-throughput suggestion engine rather than a partner in understanding.

The third barrier is the gap between computational design and experimental realisation. Generative inverse design may propose structures that satisfy formal objectives but are difficult to synthesise, unstable under realistic conditions, or incompatible with available processing routes. Chen, Zhang, Nie, Li, and Pan's review of generative models for inorganic solid materials highlights both the promise of inverse design and the need for stronger validity, stability, and synthesizability constraints [29]. This gap explains why physics-informed modelling and autonomous experimentation must be connected to generative AI rather than added later.

The fourth barrier is integration across software, hardware, institutions, and scientific culture. Self-driving laboratories require interoperable instruments, reliable robotics, standardised data capture, safety protocols, and researchers willing to trust adaptive experimental systems. Tom and colleagues emphasised that autonomous laboratories must solve practical issues of orchestration, reproducibility, and workflow design, not merely algorithmic optimisation [14]. The cultural challenge is equally significant: scientists must learn to treat AI outputs neither as unquestionable authority nor as opaque automation, but as evidence-bearing contributions to a shared discovery process.

Future Directions and Roadmap

The future of AI-driven materials science will likely depend on foundation models that can integrate text, structure, spectra, images, simulation outputs, and experimental histories. Literature-based learning has already shown that hidden knowledge can be extracted from scientific text, suggesting that future systems may use papers, databases, and laboratory records as connected sources of hypothesis generation [8]. Graph neural networks and atomistic line graph neural networks provide complementary structure-aware foundations for reasoning about molecules and crystals [15, 30]. A mature co-scientist will need to combine these modalities rather than optimise within only one representation.

A second direction is the development of closed-loop discovery infrastructures that are modular, auditable, and transferable across laboratories. Self-driving laboratories should evolve from impressive individual demonstrations into standardised architectures linking design algorithms, robotic platforms, characterisation tools, and data infrastructure. The thin-film autonomous platform developed by MacLeod and colleagues and the broader self-driving laboratory roadmap

synthesised by Tom and colleagues point toward this integrated future [6, 14]. The next challenge is to make such systems reproducible beyond elite laboratories with highly specialised hardware and software teams.

A third direction is deeper integration of generative design with physics-informed validation and experimental feedback. Generative systems should not only propose candidates but also estimate uncertainty, check physical constraints, predict synthesizability, and revise hypotheses after failed experiments. Physics-informed neural networks and hybrid atomistic models provide the grounding needed to make generative AI more reliable, while active learning supplies a mechanism for deciding which uncertainties should be resolved experimentally [10, 12, 18]. This convergence could turn inverse design from a candidate-generation exercise into a disciplined cycle of scientific reasoning.

A final direction concerns governance, ethics, and intellectual responsibility in AI-assisted discovery. As AI systems contribute to hypotheses, experimental choices, and interpretation, materials science will need clearer norms for attribution, reproducibility, intellectual property, safety, and accountability. The co-scientist model should therefore include not only technical architectures but also transparent records of model provenance, uncertainty, human oversight, and experimental decision pathways. Fully autonomous discovery will be scientifically credible only if it is also interpretable, auditable, and aligned with human-defined research values.

Conclusion

Artificial intelligence in materials science has evolved from a pattern-recognition tool into a broader discovery infrastructure. Predictive models accelerated screening, generative systems expanded the space of possible materials, physics-informed methods improved scientific grounding, and autonomous laboratories began to connect computation with real experimental action.

The co-scientist paradigm captures this transformation by describing AI as an active collaborator in hypothesis generation, design, experimentation, and interpretation. This does not mean that AI replaces human scientists, but that discovery is increasingly organised around a partnership between human judgement and machine intelligence.

Realising this paradigm will require better data, more interpretable models, stronger links between design and synthesis, and trustworthy autonomous laboratory architectures. The future of materials science will be shaped not only by what AI can predict, but by how responsibly humans and AI learn to discover together.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 26 Jul 2025

Revised: 13 Sep 2025

Accepted: 28 Oct 2025

Published online: 18 January 2026

Rights and permissions

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- de Pablo JJ, Jackson NE, Webb MA, Chen LQ, Moore JE, Morgan D, et al. New frontiers for the materials genome initiative. *npj Comput Mater.* 2019;5(1):41.
- Ramprasad R, Batra R, Pilania G, Mannodi-Kanakthodi A, Kim C. Machine learning in materials informatics: Recent applications and prospects. *npj Comput Mater.* 2017;3(1):54.
- Butler KT, Davies DW, Cartwright H, Isayev O, Walsh A. Machine learning for molecular and materials science. *Nature.* 2018;559(7715):547-55.
- Sanchez-Lengeling B, Aspuru-Guzik A. Inverse molecular design using machine learning: Generative models for matter engineering. *Science.* 2018;361(6400):360-5.
- Häse F, Roch LM, Aspuru-Guzik A. Next-generation experimentation with self-driving laboratories. *Trends Chem.* 2019;1(3):282-91.
- MacLeod BP, Parlange FG, Morrissey TD, Häse F, Roch LM, Dettelbach KE, et al. Self-driving laboratory for accelerated discovery of thin-film materials. *Sci Adv.* 2020;6(20):eaaz8867.
- Xie T, Grossman JC. Crystal graph convolutional neural networks for an accurate and interpretable prediction of material properties. *Phys Rev Lett.* 2018;120(14):145301.
- Tshitoyan V, Dagdelen J, Weston L, Dunn A, Rong Z, Kononova O, et al. Unsupervised word embeddings capture latent knowledge from materials science literature. *Nature.* 2019;571(7763):95-8.
- Long T, Fortunato NM, Opahle I, Zhang Y, Samathrakris I, Shen C, et al. Constrained crystals deep convolutional generative adversarial network for the inverse design of crystal structures. *npj Comput Mater.* 2021;7(1):66.
- Karniadakis GE, Kevrekidis IG, Lu L, Perdikaris P, Wang S, Yang L. Physics-informed machine learning. *Nat Rev Phys.* 2021;3(6):422-40.
- Jha D, Ward L, Paul A, Liao WK, Choudhary A, Wolverton C, et al. ElemNet: Deep learning the chemistry of materials from only elemental composition. *Sci Rep.* 2018;8(1):17593.
- Pun GP, Batra R, Ramprasad R, Mishin Y. Physically informed artificial neural networks for atomistic modeling of materials. *Nat Commun.* 2019;10(1):2339.
- Schmidt J, Marques MR, Botti S, Marques MA. Recent advances and applications of machine learning in solid-state materials science. *npj Comput Mater.* 2019;5(1):83.

Tom G, Schmid SP, Baird SG, Cao Y, Darvish K, Hao H, et al. Self-driving laboratories for chemistry and materials science. *Chem Rev.* 2024;124(16):9633-732.

Chen C, Ye W, Zuo Y, Zheng C, Ong SP. Graph networks as a universal machine learning framework for molecules and crystals. *Chem Mater.* 2019;31(9):3564-72.

Goodall RE, Lee AA. Predicting materials properties without crystal structure: Deep representation learning from stoichiometry. *Nat Commun.* 2020;11(1):6280.

Noh J, Kim J, Stein HS, Sanchez-Lengeling B, Gregoire JM, Aspuru-Guzik A, et al. Inverse design of solid-state materials via a continuous representation. *Matter.* 2019;1(5):1370-84.

Raissi M, Perdikaris P, Karniadakis GE. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J Comput Phys.* 2019;378:686-707.

Burger B, Maffettone PM, Gusev VV, Aitchison CM, Bai Y, Wang X, et al. A mobile robotic chemist. *Nature.* 2020;583(7815):237-41.

Batra R, Song L, Ramprasad R. Emerging materials intelligence ecosystems propelled by machine learning. *Nat Rev Mater.* 2021;6(8):655-78.

Wang J, Wang Y, Chen Y. Inverse design of materials by machine learning. *Materials (Basel).* 2022;15(5):1811.

Menon D, Ranganathan R. A generative approach to materials discovery, design, and optimization. *ACS Omega.* 2022;7(30):25958-73.

Park H, Li Z, Walsh A. Has generative artificial intelligence solved inverse materials design? *Matter.* 2024;7(7):2355-67.

Michaloglou A, Papadimitriou I, Gialampoukidis I, Vrochidis S, Kompatsiaris I. Physics-informed neural networks in materials modeling and design: A review. *Arch Comput Methods Eng.* 2025:1-38.

Shields BJ, Stevens J, Li J, Parasram M, Damani F, Alvarado JI, et al. Bayesian reaction optimization as a tool for chemical synthesis. *Nature.* 2021;590(7844):89-96.

Soldatov MA, Butova VV, Pashkov D, Butakova MA, Medvedev PV, Chernov AV, et al. Self-driving laboratories for development of new functional materials and optimizing known reactions. *Nanomaterials (Basel).* 2021;11(3):619.

Zhong X, Gallagher B, Liu S, Kaikhura B, Hiszpanski A, Han TY. Explainable machine learning in materials science. *npj Comput Mater.* 2022;8(1):204.

Pilania G. Machine learning in materials science: From explainable predictions to autonomous design. *Comput Mater Sci.* 2021;193:110360.

Chen L, Zhang W, Nie Z, Li S, Pan F. Generative models for inverse design of inorganic solid materials. *J Mater Inform.* 2021;1(1):7.

Choudhary K, DeCost B. Atomistic line graph neural network for improved materials property predictions. *npj Comput Mater.* 2021;7(1):185.