REVIEW Open access

Generative Models for Materials Science — Conceptual Capabilities and Scientific Limits: A Review Study

Laura Meier[1], Daniel Fischer[1*], Thomas Braun[2], Stefan Koch[2]

Abstract

Generative models have emerged as transformative tools in materials science, enabling the inverse design of novel materials with tailored properties by learning from vast datasets of structures and compositions. This review synthesizes recent advancements in generative approaches, including variational autoencoders, generative adversarial networks, diffusion models, and large language models. It highlights their conceptual capabilities for accelerating discovery while addressing scientific limits such as data scarcity, synthesizability, and interpretability. By examining applications in inorganic crystals, organic molecules, and energy materials, we delineate how these models bridge computational efficiency with experimental validation, yet face challenges in generalizability and physical fidelity. Future directions emphasize hybrid physics-informed architectures and closed-loop automation to overcome current barriers and unlock sustainable materials innovation.

Keywords Generative models, Materials discovery, Machine learning, Inverse design, Synthesizability, Physics-informed AI

*Correspondence:
Daniel Fischer
daniel.fischer@gmail.com
[1] Department of Materials Science and Intelligent Systems, Faculty of Engineering, University of Freiburg, Freiburg, Germany
[2] Department of AI-Driven Materials Engineering, Faculty of Engineering, Karlsruhe Institute of Technology, Karlsruhe, Germany

Introduction

Materials science stands at the intersection of physics, chemistry, and engineering, driving innovations in energy storage, electronics, catalysis, and biomedicine through the development of functional materials [1, 2]. At its core lies the pursuit of structure–property relationships: understanding how atomic arrangements, compositional variations, and processing pathways collectively determine performance. Historically, this pursuit has unfolded through iterative cycles of hypothesis formulation, experimental synthesis, and characterization. While this empirically grounded paradigm has yielded transformative technologies—from semiconductor devices to biomaterials—it remains constrained by its reliance on incremental exploration within bounded chemical subspaces [3].

The scale of the design landscape presents a profound epistemic and logistical challenge. Chemical compound space is effectively astronomical, with estimates exceeding 10^60 feasible organic molecules alone, not accounting for inorganic crystals, hybrid frameworks, or metastable phases [4]. Even with high-throughput experimentation and computational screening, only a minute fraction of this space has been interrogated. Consequently, discovery pipelines often gravitate toward well-studied compositional families, reinforcing path dependencies and limiting exposure to unconventional materials chemistries. This combinatorial explosion has catalyzed a paradigm shift from forward discovery—predicting the properties of known structures—toward inverse design, in which desired functional targets define the starting point for candidate generation [5, 6].

Artificial intelligence (AI), and specifically generative modeling, has emerged as a critical enabler of this transition. Generative models constitute a subset of deep learning systems designed not merely to classify or predict but to learn the underlying probability distributions of training data and sample from them to produce novel, statistically coherent outputs [7]. In materials science, these outputs correspond to candidate compositions, crystal structures, molecular graphs, or microstructural configurations. By approximating the latent manifolds that encode structure–property interdependencies, generative systems offer a mechanism to traverse high-dimensional design spaces with a degree of intentionality unattainable through brute-force enumeration [8].

Since 2020, architectural innovations have significantly expanded the functional scope of generative AI in materials contexts. Variational autoencoders (VAEs) have enabled continuous latent representations of materials, facilitating interpolation and property-conditioned sampling. Generative adversarial networks (GANs) have demonstrated the ability to synthesize realistic microstructures and compositional distributions through adversarial training. More recently, diffusion and score-based models have introduced iterative denoising frameworks capable of generating structurally valid crystals and molecules with improved stability and diversity profiles [9, 10]. Collectively, these architectures have shifted generative modeling from proof-of-concept demonstrations toward scalable discovery infrastructures.

A key driver of this progress has been the proliferation of large, curated datasets derived from high-throughput density functional theory (DFT) calculations, automated synthesis platforms, and open materials repositories [11, 12]. These datasets encode thermodynamic stability, electronic properties, defect energetics, and processing metadata, providing the statistical substrate upon which generative systems learn design priors. When coupled with graph-based representations, symmetry-aware encodings, and physics-informed constraints, generative models can embed domain knowledge directly into sampling processes, thereby narrowing the gap between mathematical plausibility and physical realizability.

Beyond structural generation, generative AI increasingly operates as an integrative design partner within closed-loop discovery ecosystems. Candidate materials generated by generative systems can be evaluated via simulation pipelines, filtered using synthesizability metrics, and iteratively refined through active learning. This recursive coupling of generation, evaluation, and feedback reconfigures materials discovery as a self-optimizing epistemic cycle rather than a linear screening workflow. Such integration underscores the role of generative models not simply as creative engines but as navigational instruments within complex design terrains [13].

Despite these advances, significant scientific and infrastructural limitations remain. Training datasets often reflect historical research biases, privileging stable, easily synthesizable, or industrially relevant chemistries while underrepresenting metastable or exploratory domains. Generative models trained on such data may reproduce these skews, constraining novelty and reinforcing inherited blind spots. Computational intensity constitutes another barrier: diffusion models and large latent architectures demand substantial training resources, limiting accessibility. Moreover, the translation gap between predicted candidates and experimentally realizable materials persists, shaped by factors such as kinetic barriers, synthesis conditions, and scale-up feasibility that remain only partially encoded in current generative frameworks [14, 15].

Conceptual challenges also emerge at the representation level. Materials can be encoded as sequences, graphs, voxelized volumes, or symmetry-aware tensors, each abstraction privileging certain structural features while attenuating others. The generative capacity of a model is therefore inseparable from the epistemic commitments embedded in its representational substrate. Conditioning strategies—whether property-guided, process-aware, or multimodal—further modulate design trajectories, raising questions about controllability, constraint satisfaction, and interpretability across generated candidates.

In light of this rapidly evolving landscape, this review aims to elucidate the conceptual capabilities and systemic implications of generative models in materials science. We examine how these systems navigate high-dimensional chemical spaces, operationalize inverse design, and incorporate thermodynamic, symmetry, and synthesizability

constraints. Concurrently, we critically assess scientific limits, including data bias, validation bottlenecks, computational scalability, and the interpretive opacity of latent design spaces [13–15].

Focusing on peer-reviewed literature published between 2020 and 2025, the review is organized thematically to reflect the field's layered maturation. We begin with foundational generative architectures and their domain adaptations, followed by an analysis of materials representations and conditioning mechanisms. Subsequent sections survey applications across molecular, crystalline, and microstructural material classes, highlighting domain-specific design logics. We then examine validation frameworks and translational barriers before concluding with forward-looking perspectives on hybrid generative–simulation ecosystems, multimodal data fusion, and autonomous discovery platforms [16].

Through this synthesis, the review seeks to provide researchers with a structured interpretive map of generative AI's role in materials innovation. Particular emphasis is placed on underexplored frontiers, including multimodal generative conditioning, uncertainty-aware sampling, and human-AI co-design interfaces. Ultimately, generative AI holds the potential not only to accelerate materials discovery but to democratize access to design capabilities, redistributing innovation capacity across institutional and geographic boundaries while reshaping the epistemic foundations of how materials are conceived, generated, and realized [17, 18].

Main Text

Foundations of generative models in materials science

Generative models operate by capturing the underlying distribution of training data and sampling from it to produce new instances [19]. In a materials context, this involves learning from databases of crystal structures, molecular graphs, or property profiles to generate candidates with targeted functionalities [20]. Key architectures include VAEs, which encode structures into latent spaces for efficient sampling; GANs, which use adversarial training to produce realistic outputs; and diffusion models, which iteratively denoise random noise into structured data [21, 22].

VAEs have been pivotal for their probabilistic nature, enabling uncertainty quantification in predictions [23]. For instance, they compress high-dimensional material representations into continuous latent vectors, allowing interpolation between known compounds to discover intermediates [24]. GANs excel in generating diverse samples but suffer from mode collapse, where output variety is limited [25]. Recent enhancements incorporate physics constraints, such as symmetry preservation in crystal generation, to mitigate this [26].

Diffusion models represent a breakthrough, modeling the forward diffusion of data into noise and reversing it for generation [27]. Their stability during training makes them suitable for complex systems such as proteins or alloys, where sequential refinement aligns with physical assembly processes [28]. Large language models (LLMs) adapted to specific domains treat structures as sequences (e.g., SMILES strings), enabling text-conditioned generation [29]. These foundations underscore generative models' capability to surpass enumerative screening, though limits arise from data quality—often biased toward stable, low-energy states—and the need for domain-specific adaptations [30]. The systemic interplay among generative architectures, material representations, application domains, and validation infrastructures is summarized in **Figure 1**.

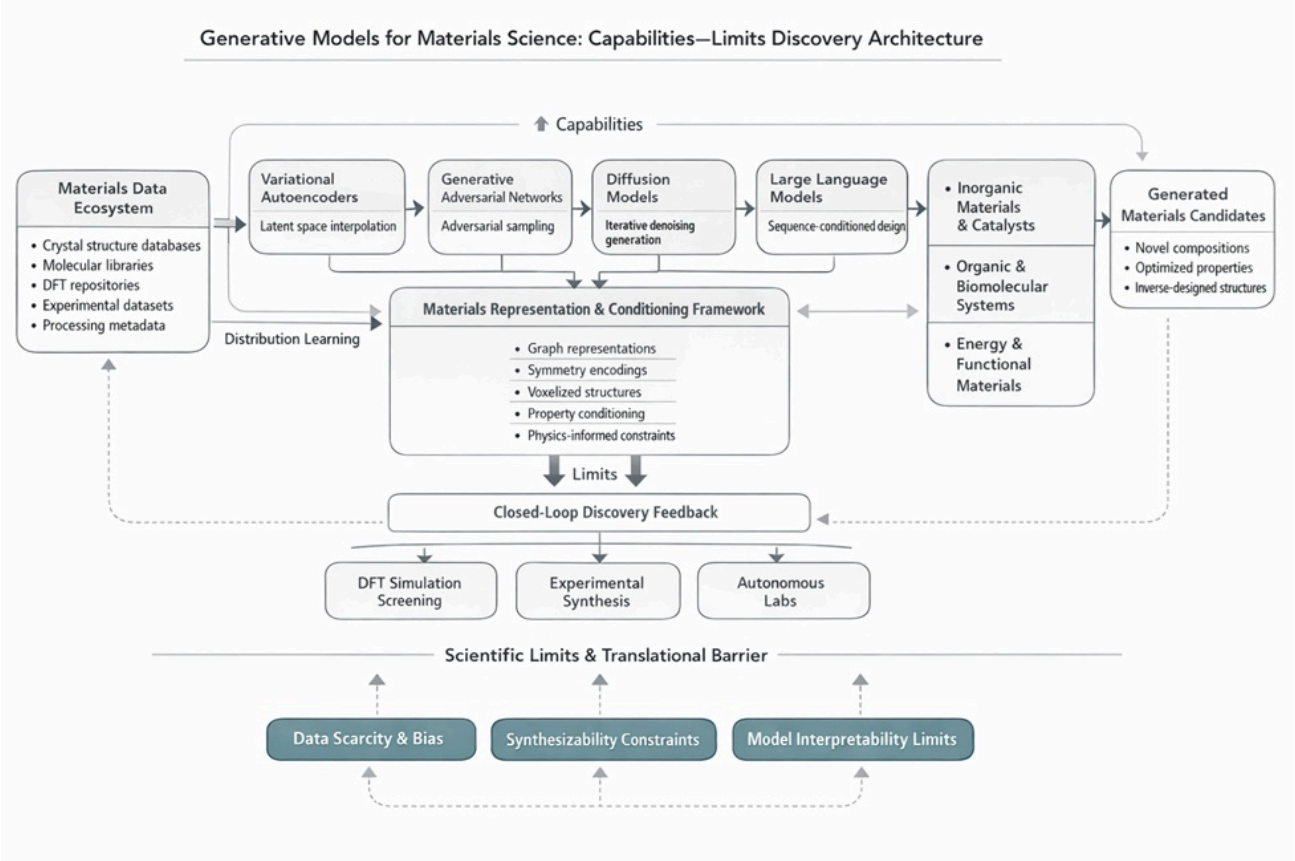

Figure 1. Generative materials discovery ecosystem: conceptual capabilities and scientific limits

A comparative synthesis of generative architectures, their conceptual capabilities, materials applications, and scientific constraints is provided in **Table 1**.

Table 1. Generative model classes in materials science: capabilities, applications, and scientific limits

Generative model	Core mechanism	Conceptual capabilities	Materials applications	Scientific limits
Variational autoencoders (VAEs)	Latent probabilistic encoding and decoding	Continuous latent interpolation; inverse design; uncertainty-aware sampling	Catalysts, polymers, and molecular design	Latent space oversmoothing; limited structural fidelity
Generative adversarial networks (GANs)	Generator–discriminator adversarial training	High-diversity sampling; microstructure synthesis; compositional exploration	High-entropy alloys; electrolytes; microstructures	Mode collapse; training instability
Diffusion models	Iterative noise denoising processes	High structural validity; multi-objective conditioning; stable training	Perovskites; proteins; semiconductors	High computational cost; slow sampling

Large language models (LLMs)	Sequence-based generative prediction	Text-conditioned design; structural grammar learning; multimodal prompts	MOFs; drug molecules; retrosynthesis	Hallucinated structures; weak physical grounding
Physics-informed generative models	Embedded thermodynamic and symmetry constraints	Physically consistent generation; improved synthesizability priors	Crystals; energy materials; functional oxides	Integration complexity; limited datasets

Material representations and conditioning strategies

Effective generative modeling hinges on representations that capture atomic interactions, periodicity, and symmetries [1, 3]. Graph-based representations, where atoms are nodes and bonds edges, are prevalent for molecules and crystals, allowing graph neural networks to encode local environments [2, 5]. Voxel grids discretize 3D space for diffusion models, facilitating the prediction of volumetric properties [4, 6].

Conditioning injects priors into the generation process, steering outputs toward desired properties such as bandgap or stability [7, 9]. Property-conditioned VAEs append target vectors to latent spaces, while GANs use auxiliary classifiers [8, 10]. Diffusion models incorporate guidance gradients and optimize multiple objectives simultaneously [11, 13]. LLMs leverage prompt engineering to convert property specifications into textual inputs [12, 14].

These strategies enhance conceptual capabilities by enabling multi-objective optimization. Yet, scientific limits include overconditioning, which reduces diversity, and representation biases, such as the omission of defects in ideal crystal models [15, 17]. Physics-informed representations, integrating DFT energies or force fields, address this by grounding generations in thermodynamics [16, 18].

Applications in inorganic materials discovery

Inorganic materials—encompassing crystalline solids, intermetallics, ceramics, and complex alloys—represent one of the most mature application domains for generative modeling in materials science. Their well-defined lattice symmetries, periodic boundary conditions, and comparatively structured thermodynamic descriptors render them particularly amenable to data-driven generation frameworks. Generative models leverage these regularities to predict stable or metastable phases, often operating within crystallographic constraint spaces that encode symmetry operations, stoichiometric bounds, and coordination environments [19, 21].

Diffusion-based generative approaches have recently demonstrated strong performance in the design of photovoltaic perovskites. By iteratively denoising candidate crystal representations, these systems can generate structurally valid ABX_3 frameworks optimized for bandgap alignment, defect tolerance, and charge-carrier mobility. Conditioning mechanisms enable simultaneous optimization of electronic properties while embedding thermodynamic stability filters derived from formation energy landscapes [20, 22]. This multi-objective generative capacity is particularly significant for halide and oxide perovskites, where compositional substitution spaces are vast yet structurally constrained.

Generative adversarial networks (GANs) have been deployed to explore compositional design spaces in high-entropy alloys (HEAs). Through adversarial training between generator and discriminator networks, GANs learn compositional distributions associated with favorable mechanical properties such as hardness, yield strength, and fracture resistance. By sampling beyond experimentally documented alloy families, these models have identified previously unexamined multi-principal element combinations with promising mechanical resilience and thermal stability [23, 25]. Such approaches reconfigure alloy discovery from incremental compositional tuning to probabilistic exploration of high-dimensional mixture spaces.

Variational autoencoders (VAEs) further enable inverse surface design in heterogeneous catalysis. Encoding catalyst surfaces as latent vectors enables property-conditioned decoding of adsorption environments tuned to specific reactant

binding energies. This facilitates the generation of catalytic interfaces optimized for selectivity, activation barriers, or resistance to poisoning across transition-metal systems [24, 26]. Importantly, latent interpolation enables the exploration of intermediate catalytic states that are not explicitly present in training datasets, expanding the accessible reaction design pathways.

Large language models (LLMs), when fine-tuned on crystallographic databases and structural descriptors, have also entered inorganic generative workflows. These models can propose metal–organic frameworks (MOFs) and coordination networks for gas storage, carbon capture, and separation applications. By learning textual–structural correspondences embedded in crystallographic records, LLMs translate functional prompts—such as pore size or adsorption selectivity—into plausible framework topologies [27, 29].

Collectively, these applications demonstrate substantial acceleration in the discovery of inorganic materials. Generative systems can reduce candidate pools from millions of enumerated structures to a few hundred high-likelihood targets, dramatically compressing computational screening burdens. Nevertheless, synthesizability remains a persistent constraint. Many generated structures occupy thermodynamically shallow minima or require kinetically inaccessible synthesis pathways. Omitting reaction kinetics, precursor availability, and processing conditions often yields candidates that are mathematically stable yet experimentally unattainable [28, 30].

Organic and biomolecular design

Generative modeling has had an equally transformative impact on the design of organic molecules and biomolecular systems, where chemical diversity and conformational flexibility far exceed those of inorganic solids. Organic chemical space—spanning small molecules, macromolecules, and bioactive compounds—presents an ideal environment for probabilistic generative exploration due to its combinatorial richness and graph-structured representations [1, 3].

Autoregressive sequence and graph-based models have been widely applied to generate drug-like compounds conditioned on pharmacological targets. By learning tokenized molecular grammars, or SMILES representations, these systems generate syntactically valid, pharmacophore-consistent molecules with predicted bioactivity profiles. Comparative studies suggest that generative pipelines can outperform traditional virtual screening by identifying viable lead compounds outside enumerated chemical libraries [2, 4]. This expands early-stage drug discovery into previously uncharted chemical territories.

Diffusion models have also emerged as powerful tools in protein and enzyme design. Operating on three-dimensional structural fields or residue contact maps, these architectures generate novel protein folds with functional binding pockets or catalytic geometries. Applications in enzyme engineering include the generation of scaffolds for substrate-specific catalysis, metabolic pathway optimization, and biosensing platforms [5, 7]. The ability to sample structurally coherent yet evolutionarily distinct folds signals a major advance in de novo biomolecular engineering.

In polymer science, variational autoencoders enable inverse design of macromolecular architectures tailored to targeted thermal, mechanical, or transport properties. By encoding polymer repeat units, branching motifs, and chain conformations, VAEs enable the optimization of properties such as thermal conductivity, elasticity, or dielectric response [6, 8]. These generative explorations are particularly valuable for energy storage membranes, thermoelectric polymers, and flexible electronics substrates.

Despite these advances, validation challenges are pronounced. Organic and biomolecular systems exhibit high conformational variability, and generative outputs may correspond to energetically unfavorable or kinetically unstable states. Models can “hallucinate” chemically implausible bonding configurations or sterically incompatible geometries when extrapolating beyond training distributions [9, 11]. To mitigate this, hybrid validation frameworks combine generative outputs with molecular dynamics simulations, quantum chemical calculations, and docking analyses. These integrative

pipelines filter generated molecules for conformational stability, reaction feasibility, and functional persistence before experimental synthesis [10, 12].

Energy and functional materials

Generative AI has become a critical design instrument in the development of energy and functional materials, where performance optimization often requires simultaneous tuning of electrochemical, electronic, and structural parameters. Battery materials represent a primary focus area, particularly solid electrolytes and electrode materials that require high ionic conductivity, electrochemical stability, and manufacturability [13, 15].

GAN-based frameworks have been applied to electrolyte discovery, yielding candidate compositions with expanded electrochemical stability windows and improved ion-transport properties. By learning structure–conductivity correlations from computational datasets, these systems propose novel ceramic and polymer electrolytes that balance conductivity with dendrite suppression and thermal resilience [14, 16].

In semiconductor research, conditioned diffusion models enable the synthesis of crystalline materials with targeted bandgaps and carrier-transport properties. Such models are especially valuable in optoelectronic design, where precise band alignment is required for applications in photodetectors, LEDs, and solar absorbers. Conditioning variables—such as lattice symmetry or orbital character—guide generative sampling toward electronically functional candidates [17, 19].

However, functional deployment introduces multi-scale modeling challenges. Many performance metrics—thermal transport, fracture resistance, degradation pathways—arise from mesoscale or macroscale phenomena that are not fully captured in atomistic training data. Current generative systems remain predominantly anchored at the microscale, limiting predictive transferability to device-level performance contexts [18, 20].

Emerging physics-informed LLMs and hybrid generative architectures attempt to bridge this divide by embedding continuum mechanics, transport equations, and degradation models into generative reasoning frameworks. Such systems aim to align atomic-scale design with engineering-scale functionality, enabling more holistic materials optimization strategies [21, 23].

Challenges and validation frameworks

Despite rapid advances in architectural and application design, generative materials modeling faces systemic constraints that limit its translational impact. Chief among these is data scarcity. High-fidelity materials datasets—particularly experimental datasets—remain limited in size and diversity. Small-data regimes increase susceptibility to overfitting, latent space collapse, and reduced generative diversity [24, 26].

Synthesizability represents a second major barrier. Generative systems frequently produce thermodynamically unstable or kinetically inaccessible candidates due to incomplete encoding of reaction pathways, precursor chemistries, or fabrication constraints. This disconnect between computational plausibility and laboratory feasibility constrains downstream adoption [25, 27].

Interpretability further complicates validation. Many generative models operate as black-box systems, obscuring the structural logic underlying candidate generation. Without mechanistic transparency, researchers face difficulty assessing why specific materials are proposed or how latent variables correspond to physical phenomena [28, 30].

To address these limitations, multi-tier validation ecosystems have emerged. Density functional theory (DFT) calculations serve as a first-line computational filter, evaluating formation energies, electronic structures, and phonon stability. Promising candidates then progress to experimental synthesis and characterization workflows for empirical validation [1, 3].

Closed-loop discovery systems represent the most advanced validation paradigm. In these infrastructures, generative models feed predictions directly into autonomous laboratories equipped with robotic synthesis, real-time characterization, and adaptive feedback algorithms. Experimental outcomes recursively refine generative priors, forming self-improving discovery cycles. While these systems mark a significant leap in capability, they require robust uncertainty handling, anomaly detection, and failure-aware learning protocols to prevent error propagation across iterative loops [2, 4].

Results and Discussion

The integration of generative models into materials science has revolutionized the field by shifting from forward prediction to inverse design. Yet, this progress is tempered by inherent limitations that must be addressed for broader adoption [1, 3, 5]. This section expands on key challenges and opportunities, organized under subheadings that delve into data-related issues, model interpretability, ethical considerations, and interdisciplinary synergies.

Data scarcity and quality: Barriers to robust generalization

A primary scientific limit of generative models lies in their dependence on high-quality, diverse datasets, which are often scarce in materials science [2, 4, 6]. Unlike image or text domains with billions of samples, materials databases like the Materials Project or ICSD contain only tens to hundreds of thousands of entries, predominantly for stable compounds [7, 9]. This scarcity leads to overfitting, where models excel on seen data but fail to generalize to novel chemistries [8, 10]. For instance, generative models trained on organic molecule datasets may overlook rare-earth elements or extreme conditions, leading to biased outputs [11, 13].

Quality issues exacerbate this: computational data from DFT approximations introduce systematic errors, while experimental data suffer from inconsistencies in measurement protocols [12, 14]. Active learning strategies that query uncertain regions for new data can mitigate scarcity but require integration with high-throughput experiments [15, 17]. Future efforts should prioritize curated, multimodal datasets combining structures, spectra, and synthesis routes to enhance model robustness [16, 18].

Interpretability and physical fidelity: Bridging the black-box gap

Generative models' conceptual strength in producing diverse candidates is undermined by their opacity, making it difficult to discern why a structure is proposed [19, 21]. In materials science, where physical laws govern viability, lack of interpretability hinders validation and trust [20, 22]. Techniques like attention mechanisms in LLMs or saliency maps in GANs provide partial insights, highlighting influential atoms or bonds [23, 25]. However, these fall short for complex phenomena such as superconductivity, where emergent behaviors defy simple explanations [24, 26].

Physics-informed generative models address this by embedding conservation laws or symmetries into their architectures, ensuring that their outputs respect thermodynamics [27, 29]. For example, equivariant diffusion models preserve rotational invariance in crystal generation, improving fidelity [28, 30]. Despite these advances, limits persist in multi-scale phenomena, such as defect propagation, necessitating hybrid approaches that couple generative outputs with ab initio simulations for post-hoc verification [1, 3].

Synthesizability and experimental validation: From virtual to real

A critical limit is the disconnect between generated candidates and laboratory realization, as models often prioritize property optimization over synthetic feasibility [2, 4]. Generative outputs may propose unstable intermediates or require impractical conditions, ignoring kinetic barriers [5, 7]. Metrics like synthetic accessibility scores (SAS) integrated into conditioning help, but rely on heuristic approximations [6, 8].

Closed-loop automation, where generative models interface with robotic synthesis platforms, represents a transformative capability [9, 11]. Recent demonstrations in alloy design showcase iterative refinement: models propose, experiments

validate, and feedback retrains [10, 12]. Yet, challenges include scalability—high-cost experiments limit iterations—and handling failures, such as toxic byproducts [13, 15]. Expanded validation frameworks that incorporate machine-learning surrogates for rapid screening are essential to bridge this gap [14, 16].

Ethical and societal implications: Responsible AI in materials innovation

Beyond technical limits, generative models raise ethical concerns in materials science, particularly for dual-use applications like advanced energetics [17, 19]. Democratized access could accelerate beneficial discoveries, such as sustainable batteries, but also risks misuse without safeguards [18, 20]. Bias in training data—often skewed toward the priorities of industrialized nations—may perpetuate inequities in global materials research [21, 23].

Opportunities lie in transparent governance: open-source models with built-in constraints for sensitive domains [22, 24]. Interdisciplinary collaboration with ethicists can embed fairness metrics to ensure diverse representation in datasets [25, 27]. As generative AI matures, balancing innovation with responsibility will define its societal impact [26, 28].

Interdisciplinary synergies: Integrating with emerging technologies

Generative models' capabilities extend through synergies with quantum computing, robotics, and big data analytics [1, 29]. Quantum-enhanced sampling could address the exponential complexity of electronic structure, while AI agents can automate end-to-end workflows [2, 30]. These integrations promise to overcome current limits, fostering a new era of accelerated discovery [3, 5].

Conclusion

Generative models have profoundly expanded the conceptual toolkit of materials science, enabling inverse design that navigates immense chemical spaces with unprecedented efficiency. From VAEs and GANs to diffusion models and LLMs, these approaches have demonstrated applications in inorganic, organic, and energy materials, accelerating innovations in sustainability and technology. However, scientific limits—data scarcity, interpretability deficits, challenges to synthesizability, and ethical considerations—underscore the need for cautious advancement.

Looking forward, hybrid physics-informed architectures, closed-loop experimentation, and ethical frameworks will be pivotal in realizing the full potential of generative AI. By addressing these barriers, the field can transition from conceptual promise to tangible scientific breakthroughs, ultimately yielding materials that address global challenges such as climate change and resource scarcity.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 14 Sep 2022 Revised: 26 Nov 2022 Accepted: 29 Dec 2022
Published online: 18 July 2023

References

Zeni C, Pinsler R, Zügner D, Fowler A, Horton M, Fu X, et al. A generative model for inorganic materials design. Nature. 2025;639(8055):624-32.
https://doi.org/10.1038/s41586-025-08628-5.

Otyepka M, Pykal M, Otyepka M. Advancing materials discovery through artificial intelligence. Appl Mater Today. 2025;47:102981.
https://doi.org/10.1016/j.apmt.2025.102981.

Sun Y, Chen L, Jing Z, Li YY, Kim D, Gao JY, et al. Generative ai for the design of molecules: Advances and challenges. J Chem Inf Model. 2025;65(23):12668-90.
https://doi.org/10.1021/acs.jcim.5c02234.

Xie E, Wang X, Siepmann JI, Chen H, Snurr RQ. Generative ai for design of nanoporous materials: Review and future prospects. Digit Discov. 2025;4(9):2336-63.
https://doi.org/10.1039/D5DD00221D.

Zhang L, Liu Z, Ni B, Wang Q. Large language models (LLMS) for materials design. Adv Funct Mater. 2025;36(30).
https://doi.org/10.1002/adfm.202525897.

Bai X, Zhang X. Artificial intelligence-powered materials science. Nano Micro Lett. 2025;17(1):135.
https://doi.org/10.1007/s40820-024-01634-8.

Uddin M, Arfeen SU, Alanazi F, Hussain S, Mazhar T, Rahman MA. A critical analysis of generative ai: Challenges, opportunities, and future research directions. Arch Comput Methods Eng. 2025;33:1763-93.
https://doi.org/10.1007/s11831-025-10355-z.

Peckham O, Raines J, Bulsink E, Goudswaard M, Gopsill J, Barton D, et al. Artificial intelligence in generative design: A structured review of trends and opportunities in techniques and applications. Designs. 2025;9(4):79.
https://doi.org/10.3390/designs9040079.

Unni R, Zhou M, Wiecha PR, Zheng Y. Advancing materials science through next-generation machine learning. Curr Opin Solid State Mater Sci. 2024;30:101157.
https://doi.org/10.1016/j.cossms.2024.101157.

Menon D, Ranganathan R. A generative approach to materials discovery, design, and optimization. ACS Omega. 2022;7(30):25958-73.
https://doi.org/10.1021/acsomega.2c03264.

Pyzer-Knapp EO, Manica M, Staar P, Morin L, Ruch P, Laino T, et al. Foundation models for materials discovery – current state and future directions. npj Comput Mater. 2025;11(1):61.
https://doi.org/10.1038/s41524-025-01538-0.

Fuhr AS, Sumpter BG. Deep generative models for materials discovery and machine learning-accelerated innovation. Front Mater. 2022;9:865270.
https://doi.org/10.3389/fmats.2022.865270.

He R, Cao J, Tan T. Generative artificial intelligence: A historical perspective. Natl Sci Rev. 2025;12(5):nwaf050.

Braham EJ, Ruddock JM, Hardin JO. Generating and leveraging explanations of ai/ml models in materials and manufacturing research. Patterns. 2025;6(9):101340.
https://doi.org/10.1016/j.patter.2025.101340.

Gregores Coto A, Eike Precker C, Andersson T, Laukkanen A, Suhonen T, Rey Rodriguez P, et al. The use of generative models to speed up the discovery of materials. Comput Methods Mater Sci. 2023;23(1):13-26.
https://doi.org/10.7494/cmms.2023.1.0802.

Li H, Ding X, Zhang Z. Inverse design with deep generative models: Next step in materials discovery. Natl Sci Rev. 2022;9(8):nwac111.

Liu Y, Yang Z, Yu Z. Generative artificial intelligence and its applications in materials science: Current situation and future perspectives. J Materiomics. 2023;9(4):798-816.
https://doi.org/10.1016/j.jmat.2023.05.001.

Zhang J, Zheng Y, Wang Z. Materials discovery acceleration by using conditional generative methodology. npj Comput Mater. 2025;11(1):193.
https://doi.org/10.1038/s41524-025-01930-w.

Han Y, Buehler MJ. Generative artificial intelligence for advancing discovery and design in biomateriomics. Intell Comput. 2025;4:0117.
https://doi.org/10.34133/icomputing.0117.

Zhou H, Xu J, Qin X, Zhang J, Zou W, Shakouri M, et al. Machine learning-driven material intelligence research and development. Nano Res. 2025;18(12):94908095.
https://doi.org/10.26599/NR.2025.94908095.

Baek SJ, Lee M. Recent advances in the inverse design of silicon photonic devices using deep generative models. PeerJ Comput Sci. 2025;11:e2895.
https://doi.org/10.7717/peerj-cs.2895.

Singh AK, Kumar R. Machine learning-driven materials discovery: Unlocking next-generation functional materials – a review. Mater Discov. 2025;22:100014.
https://doi.org/10.1016/j.md.2025.100014.

Merchant A, Batzner S, Schoenholz SS, Aykol M, Cheon G, Cubuk ED. Scaling deep learning for materials discovery. Nature. 2023;624(7990):80-5.
https://doi.org/10.1038/s41586-023-06735-9.

Fuhr AS, Sumpter BG. Artificial intelligence driving materials discovery? Perspective on the article: Scaling deep learning for materials discovery. Chem Mater. 2024;36(8):3501-3.

https://doi.org/10.1021/acs.chemmater.4c00643.

Yu S, Ran N, Liu J. From large language models to ai agents in energy materials research: Enabling discovery, design, and automation. AI Agent. 2025;1:9.
https://doi.org/10.20517/aiagent.2025.03.

Leung EKH, Balan N, Lee CKH, Xie S. The Generative Artificial Intelligence large language product design multi-model framework for manufacturing operations. J Oper Res Soc. 2025:1–27.
https://doi.org/10.1080/01605682.2025.2570407.

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.
https://doi.org/10.1038/s41524-020-00352-0.

Court CJ, Yildirim B, Jain A. Accelerating materials discovery with machine learning. J Phys Mater. 2021;4(2):022001.

Han Y, Buehler MJ. Automating alloy design and discovery with physics-aware multimodal multiagent ai. Proc Natl Acad Sci U S A. 2025;122(44):e2414074122.
https://doi.org/10.1073/pnas.2414074122.

Ghafarollahi A, Buehler MJ. Automating alloy design and discovery with physics-aware multimodal multiagent ai. Proc Natl Acad Sci U S A. 2025;122(44):e2414074122.