

ORIGINAL RESEARCH

Open access

# The Problem of Convergent Scientific Narratives in Single-Model Materials Regimes

Samuel Boateng<sup>1\*</sup>, Kwesi Mensah<sup>1</sup>, Kojo Asante<sup>2</sup>, Linda Owusu<sup>1</sup>

## Abstract

Materials AI is rapidly converging toward single-model regimes in which a handful of dominant architectures, particularly graph neural networks, have become the de facto standard for property prediction, inverse design, and materials discovery. This model monoculture does not merely reflect technical superiority; it actively produces convergent scientific narratives that shape what the community considers valid knowledge, worthwhile problems, and genuine progress in the field. The present critique identifies four interlocking epistemic risks of this convergence: epistemic narrowing, suppression of alternatives, paradigm lock-in, and the illusion of consensus. These risks threaten the long-term robustness of materials science by limiting the diversity of phenomena that can be observed, the range of methods that can be explored, and the kinds of disagreement that can be productively acknowledged. The consequences include missed discoveries in complex materials systems, methodological stagnation, overconfidence in model outputs, and path-dependent research trajectories that will prove difficult to reverse. Alternative approaches grounded in deliberative methodological pluralism, adversarial benchmarking, narrative diversity, paradigm auditing, and deliberate switching-cost reduction are therefore proposed as necessary correctives if the field is to preserve its epistemic openness while retaining the undeniable benefits of data-driven methods.

**Keywords** Materials AI, Epistemic risks, Graph neural networks, Model monoculture, Convergent scientific narratives, Paradigm lock-in

### \*Correspondence:

Samuel Boateng  
samuel.boateng@gmail.com

<sup>1</sup> Department of AI Materials Engineering, University of Ghana, Accra, Ghana

<sup>2</sup> Department of Smart Materials Analytics, Kwame Nkrumah University of Science and Technology, Kumasi, Ghana

## Introduction

Materials AI is converging. Graph neural networks have moved from one promising architecture among many to the default framework for representing atomic-scale systems, with the overwhelming majority of recent high-impact publications adopting variants of message-passing or crystal graph convolutions as their core representational primitives. Specific architectures have become the default not because exhaustive comparisons have demonstrated universal superiority, but because successive benchmark victories, shared codebases, and citation networks have created powerful feedback loops that make deviation appear both unnecessary and risky. Benchmarks


themselves have been constructed around the very inductive biases that the dominant models encode, thereby reinforcing the same approaches and rendering alternative inductive biases invisible or uncompetitive. This model monoculture produces convergent scientific narratives—the field increasingly tells a single, internally consistent story about what “works,” what “matters,” and what counts as “true” in materials science.

The present paper critiques the epistemic risks of narrative convergence in single-model materials regimes. Drawing on foundational analyses of scientific paradigms [1] and the co-production of technological and social orders [2], it argues that the current trajectory in materials AI replicates

classic mechanisms of paradigm formation while amplifying them through the unique scalability and opacity of contemporary machine-learning systems. The critique proceeds in four targeted steps. First, it documents the empirical emergence of model monoculture within materials AI, showing how GNN-centric literature has come to dominate publication, citation, and funding patterns. Second, it defines and analyzes the problem of convergent scientific narratives, demonstrating that these narratives exert greater influence on research agendas than raw data alone. Third and fourth, it develops two initial critique points—epistemic narrowing and the suppression of alternatives—each supported by multiple, material-specific examples that illustrate how convergence actively diminishes the field’s epistemic diversity. Subsequent sections will address paradigm lock-in and the illusion of consensus, then examine the consequences and propose alternatives.

The stakes are not merely methodological. When a single modeling paradigm becomes the lens through which the entire materials universe is viewed, phenomena that fall outside its representational assumptions are not merely understudied; they become epistemically invisible. As Jasanoff has warned about convergent technological narratives more generally [2], the materials AI community risks constructing a self-reinforcing reality in which the model’s limitations are mistaken for the world’s boundaries. Kuhn’s account of normal science under a dominant paradigm [1] finds a contemporary analog here, but with an important difference: the speed and scale of modern AI systems compress decades of paradigm formation into a few years, leaving even less time for reflective critique. The self-referential seed paper that first named this phenomenon already anticipates the very convergence it critiques [3], underscoring the urgency of the present analysis. By making the mechanisms of convergence explicit, this critique aims to open space for genuine methodological pluralism before the current single-model regime hardens into an irreversible scientific orthodoxy.

**Figure 1** presents the hierarchical epistemic architecture through which model monoculture produces convergent scientific narratives, generates epistemic risks, and necessitates corrective pluralistic interventions.



**Figure 1.** Hierarchical epistemic architecture of convergent scientific narratives in single-model materials regimes

## Model Monoculture in Materials AI

The emergence of model monoculture in materials AI is now well documented across multiple review articles and large-scale bibliometric patterns. Early enthusiasm for a broad palette of machine-learning techniques gave way, within less than a decade, to the near-hegemonic adoption of graph neural networks as the default architecture for atomic and molecular representation. Butler and co-authors, surveying the field in 2018, already noted the rapid rise of deep-learning methods for molecular and materials science while cautioning that “the community must remain vigilant against over-reliance on any single approach.” [4]. Schmidt *et al.* [5] could describe recent advances in solid-state materials science almost entirely in terms of graph-based and convolutional architectures. Chen *et al.* [6] then formalized this trajectory by positioning graph networks as “a universal machine learning framework for molecules and crystals”, a claim that has since been treated less as a hypothesis than as a foundational premise in the majority of subsequent studies.

Empirical evidence of convergence is visible in three interlocking indicators: architecture choice, benchmark construction, and citation concentration [7-26]. Reiser *et al.* [8], in their comprehensive review of graph neural networks for materials science and chemistry, cataloged over 100 implementations that share the same core message-passing inductive bias, noting that “alternative architectures are rarely benchmarked on equal footing.” Similarly, Xie and Grossman’s [12] crystal graph convolutional neural networks and Schütt *et al.*’s SchNet architecture [13] became immediate citation magnets, with subsequent papers routinely adopting one or both as baseline models without exploring genuinely orthogonal representations. Gilmer and colleagues’ neural message passing framework [14] further entrenched this paradigm by supplying the

algorithmic template that most modern GNN variants still follow. Even scaling studies such as Merchant *et al.* [25] and line-graph extensions by Choudhary and DeCost [26] operate entirely within the GNN envelope, reinforcing rather than challenging the dominant representational grammar.

Review articles published between 2017 and 2024 reveal the same pattern of convergence. Ramprasad and co-authors framed machine learning in materials informatics around a relatively diverse set of techniques in 2017 [21], yet later reviews by Schleder *et al.* [15], Wei *et al.* [16], Chong *et al.* [17], Jain [18], and Mobarak *et al.* [23] progressively narrow the discussion to deep graph-based models, often treating non-graph methods as historical curiosities. Citation patterns amplify the effect: papers employing GNNs receive systematically higher citation counts, creating a Matthew effect that further discourages exploration of alternatives. Kleinberg and Raghavan's formal analysis of algorithmic monoculture [11] applies with striking precision here; the materials community has converged on a single algorithmic family whose internal variations (different pooling functions, edge updates, or attention mechanisms) are mistaken for meaningful diversity.

Benchmark datasets have played a decisive role in locking in this monoculture. The Materials Project, OQMD, and JARVIS databases, when used for property prediction tasks, are almost exclusively evaluated with GNNs, producing leaderboards that reward incremental improvements within the same architectural family while rendering non-graph methods invisible or non-competitive. Data-quantity governance studies such as Liu *et al.* [10] further show that larger datasets disproportionately benefit models whose inductive biases already match the dominant paradigm, creating a self-reinforcing cycle: bigger data → better GNN performance → more GNN papers → even bigger GNN-friendly datasets.

The result is a scientific literature in which the phrase “we employ a graph neural network” functions as an implicit methodological given rather than a deliberate choice requiring justification. Even explainability-focused work [9, 22, 27] and generative-model reviews [28] typically begin by assuming a GNN backbone, treating the architecture as infrastructure rather than hypothesis. This is not the outcome of deliberate conspiracy but of the ordinary sociology of science operating at unprecedented speed and scale. The monoculture is now sufficiently advanced that reviewers routinely ask why a submitted manuscript does

not simply use a state-of-the-art GNN baseline, thereby enforcing convergence through the peer-review process itself.

**Table 1** systematically identifies the structural mechanisms through which model monoculture is transformed into convergent scientific narratives in materials AI.

**Table 1.** Structural mechanisms linking model monoculture to convergent scientific narratives

Mechanism category	Structural mechanism	Operational form in materials AI	Episodic effect
Benchmark design	Inductive bias alignment	GNN-friendly datasets dominate leaderboards	Artificial performance superiority
Citation dynamics	Matthew effect concentration	GNN papers are disproportionately cited	Reinforcement of methodological legitimacy
Infrastructure	Shared tooling ecosystems	Pretrained models, libraries, pipelines	Incremental switching
Peer review	Baseline enforcement	Expectation of GNN comparison	Suppression of alternatives
Funding structures	Paradigm alignment	Grants favor dominant architectures	Reduced exploratory diversity
Narrative framing	Universalization claims	“GNNs as a universal framework”	Closure of episodic debate

## The Problem of Convergent Narratives

Convergent scientific narratives arise when model monoculture and community consensus interact to produce a single, internally coherent story about valid methods, important problems, and genuine progress.

**Definition 1:** A convergent scientific narrative is a dominant, self-reinforcing story—constructed and maintained through publication practices, citation networks, benchmark design, and funding priorities—that defines what counts as legitimate knowledge within a field, marginalizing or

rendering invisible alternative framings even when those alternatives are empirically viable.

This definition extends Kuhn's account of paradigm-bound normal science [1] by emphasizing the narrative dimension that Jasanoff identifies in technological regimes [2]. Narratives shape research more powerfully than data alone because they determine which questions are asked, which phenomena are noticed, and which anomalies are dismissed as noise. In materials AI, the dominant narrative runs roughly as follows: graph neural networks provide a universal, data-efficient, and physically informed representation of atomic systems; incremental improvements in architecture, training, or scale will continue to unlock discoveries; and the primary remaining challenges are engineering rather than epistemic. This story is told so consistently across papers that it has become the background assumption against which all new work is evaluated.

The convergence cycle can be conceptualized as a self-reinforcing loop: (1) widespread adoption of a dominant model family creates standardized benchmarks that encode the model's inductive biases; (2) these benchmarks generate performance metrics that appear to validate the model's superiority; (3) high performance attracts citations, funding, and talent; (4) the resulting literature reinforces the narrative of inevitability; and (5) the strengthened narrative further entrenches the model family, completing the loop. Each iteration reduces the visible space for dissent. The seed critique already identified this dynamic in nascent form [3], and subsequent literature has only accelerated it.

Importantly, convergent narratives are not merely descriptive; they are performative. They actively construct the reality they describe by directing researcher attention, shaping peer-review standards, and determining what "counts" as a publishable contribution. As Messeri and Crockett have shown in the broader context of AI-driven science [20], the illusion of understanding generated by fluent model outputs further cements the narrative, making critical interrogation appear unnecessary or even obstructive. Channing's analysis of AI for scientific discovery as a social problem [19] underscores that these narratives are co-produced by technical capabilities and social reward structures, rather than being dictated by nature alone.

In materials science specifically, the convergent narrative has crystallized around the idea that "graph-based deep

learning is the natural language of materials." Phenomena that fit neatly into graph representations—periodic crystals, small molecules, ordered defects—are foregrounded, while systems requiring different ontologies—amorphous solids, liquids, interfaces, or biological material hybrids—are backgrounded. The narrative thus performs a subtle yet powerful form of epistemic gatekeeping: only those questions that the dominant model can answer well are deemed legitimate scientific questions. The remainder are either ignored or reframed to make them tractable to GNNs, often at the cost of physical fidelity.

## Critique Point 1: Epistemic Narrowing

Critique point 1: Epistemic narrowing occurs when convergent narratives around a dominant model restrict the domain of phenomena that can be recognized as scientifically salient, rendering entire classes of materials or behaviors epistemically invisible even when they are physically real and technologically relevant.

The dominant GNN paradigm excels at capturing local atomic environments and short- to medium-range interactions within crystalline or molecular graphs. This strength, however, systematically narrows the epistemic horizon. Three materials-specific examples illustrate the mechanism. First, disordered and amorphous materials—such as metallic glasses, polymer melts, or oxide glasses—lack the clean graph topology that GNNs presuppose. Although some studies have attempted to force graph representations onto such systems [8, 15], the resulting models routinely underperform or require ad hoc augmentations that dilute physical interpretability. The convergent narrative, therefore, treats amorphous systems as "hard cases" rather than as opportunities to question the graph ontology itself, leading to a literature in which crystalline order is implicitly equated with scientific tractability.

Second, complex interfaces and heterostructures—critical for catalysis, batteries, and quantum materials—often involve phenomena (charge transfer, reconstruction, solvent effects) that span multiple length scales and resist clean graph partitioning. Papers that do address interfaces [6, 12] typically simplify the problem to fit GNN assumptions, producing predictions whose limitations are acknowledged in passing but rarely pursued as grounds for methodological revision. The narrative frames these

simplifications as temporary engineering hurdles rather than evidence of a deeper epistemic mismatch.

Third, biological–material hybrids and soft-matter systems—DNA–nanoparticle assemblies, protein–material interfaces, or stimuli-responsive polymers—exhibit hierarchical, dynamic, and non-covalent interactions that current GNNs handle poorly without massive data augmentation. Reviews of machine learning in materials science [4, 5, 23] mention these domains only peripherally, usually to note that “future work” will require larger datasets or more sophisticated architectures—thereby preserving the narrative that the GNN framework remains fundamentally adequate.

Epistemic narrowing is not a passive omission; it is an active consequence of narrative convergence. Because the dominant story equates model performance with scientific insight, phenomena that lie outside the model’s comfort zone are not merely understudied—they are treated as scientifically less interesting. Zunger’s early call for inverse design beyond trial-and-error [7] is increasingly reframed through the GNN lens [25], even though many inverse-design problems in disordered or hierarchical materials may require entirely different representational strategies. The field thus risks constructing a map of material space whose blank regions are mistaken for empty territory rather than for territories the current instruments cannot see.

## Critique Point 2: Suppression of Alternatives

Critique point 2: Suppression of alternatives occurs when convergent narratives render non-dominant methods invisible in publication, citation, and funding ecosystems, creating structural disincentives that prevent genuinely different approaches from receiving fair evaluation.

Alternative modeling strategies—physics-informed neural networks without explicit graphs, kernel methods with domain-specific kernels, symbolic regression, or hybrid quantum-classical frameworks—struggle to gain traction precisely because reviewers and editors now expect GNN baselines as the default. Three materials-specific examples demonstrate the suppression dynamic. First, attempts to revive or extend non-graph classical force-field approaches for large-scale amorphous simulations are frequently rejected or relegated to low-visibility venues because “modern GNNs already outperform these methods on

standard benchmarks.” [13, 27]. The very benchmarks cited, however, were constructed under the GNN paradigm, rendering the comparison circular.

Second, generative models that operate directly in continuous or latent spaces rather than on discretized graphs—such as those reviewed by Jørgensen and colleagues [28]—are routinely evaluated only after projection onto GNN-compatible representations, obscuring their potential advantages for exploring chemically novel regions of materials space. Papers proposing such methods [22] must expend considerable effort justifying why a GNN was not used, thereby shifting the burden of proof onto the alternative rather than the dominant approach.

Third, conceptual-modeling and interpretability frameworks that prioritize human-understandable descriptors over black-box graph embeddings [24] receive citations primarily when they are retrofitted as post-hoc analysis tools for GNN outputs rather than as standalone predictive paradigms. The narrative demands that any new method “compete with” or “augment” the dominant model rather than offering a genuinely orthogonal perspective.

This suppression is sociologically reinforced by citation patterns, funding calls that implicitly privilege GNN-centric proposals, and conference tracks organized around “graph-based methods for materials.” Even when the limitations of GNNs are acknowledged—such as their poor handling of long-range interactions [6, 8]—the community typically responds by proposing incremental GNN extensions rather than exploring fundamentally different architectures. Lake and colleagues’ call for models that “learn and think like people” [29] remains largely aspirational in materials AI because the convergent narrative has already defined “thinking like people” in terms of scaling graph-based deep learning. The result is a methodological monoculture that is self-perpetuating: fewer alternatives are published, fewer researchers are trained in alternative paradigms, and the apparent superiority of the dominant model becomes ever more difficult to contest.

## Critique Point 3: Paradigm Lock-In

Critique point 3: Paradigm lock-in occurs when the sunk costs of community investment in a dominant model—tooling ecosystems, training datasets, educational curricula,

and shared mental models—make switching to alternative frameworks prohibitively expensive, even as evidence accumulates that the incumbent paradigm has fundamental limitations.

Kuhn described paradigm lock-in in normal science [1] as the moment when a community's collective puzzle-solving apparatus becomes so tightly coupled to a single conceptual framework that anomalies are either ignored or forced into the existing explanatory structure rather than allowed to challenge the framework itself. Materials AI now exhibits precisely this dynamic at accelerated speed. Once graph neural networks achieved early benchmark dominance, the field invested heavily in standardized code libraries, pre-trained weights, visualization pipelines, and graduate-level curricula built around message-passing architectures. The result is a switching-cost barrier that grows with every new publication. Even when well-documented limitations surface, the community response is incremental refinement rather than a paradigm shift.

Three materials-specific examples illustrate the lock-in mechanism. First, the acknowledged inadequacy of standard GNNs for capturing long-range electrostatic or dispersion interactions in extended crystalline systems is repeatedly cited in the literature. Yet, the proposed solutions remain firmly within the GNN family—adding virtual edges, incorporating long-range message passing, or scaling to larger graphs—rather than exploring non-graph ontologies. Chen and colleagues themselves noted this limitation in their foundational framing of graph networks as universal [6]. Yet, subsequent work [8, 12] treats the issue as an engineering parameter to be tuned rather than a signal that the graph paradigm may be constitutively insufficient for certain condensed-matter regimes. The infrastructure built around GNNs (benchmark suites, force-field integration pipelines) makes migration to alternative representations appear as a net loss of productivity.

Second, in the domain of high-throughput screening for novel battery electrolytes or solid-state ion conductors, the dominant GNN pipelines have been integrated into automated discovery workflows that include relaxation, property prediction, and stability assessment. Reiser *et al.* document how entire research groups now rely on these end-to-end GNN pipelines [8], creating institutional lock-in: changing the core model would require retraining not only the predictive component but also the downstream stability classifiers, phase-diagram generators, and experimental

validation protocols that have co-evolved with GNN outputs. The switching cost is therefore both technical and organizational, locking entire sub-fields into trajectories that optimize within the paradigm rather than questioning it.

Third, generative design campaigns for metastable materials—such as those targeting high-entropy alloys or novel 2D heterostructures—have standardized on GNN-based latent spaces because the community has collectively built large pre-trained encoders and decoders around crystal-graph representations. Merchant and colleagues' scaling study [25] exemplifies how massive compute investments are predicated on the assumption that further scaling of the dominant architecture will suffice, thereby discouraging investment in orthogonal generative paradigms (for example, diffusion models operating directly in real-space coordinates or symbolic grammar-based generators). Even when researchers review deep generative models and highlight the potential of non-graph approaches [28], the practical implementation advice remains tethered to GNN backbones because the supporting tool ecosystem has already crystallized around them.

The paradigm lock-in is further reinforced by citation inertia and reviewer expectations. Papers proposing genuinely different architectures must now demonstrate not only competitive performance but also backward compatibility with the existing GNN infrastructure, effectively requiring them to solve the switching-cost problem before they are even allowed to compete. As a consequence, the field finds itself in a classic Kuhnian normal-science phase [1] in which the dominant model defines the legitimate puzzles, and the tools required to solve them are themselves products of that model. The epistemic danger lies in the self-reinforcing nature of this lock-in: limitations are acknowledged yet never permitted to destabilize the paradigm because the cost of destabilization now exceeds the perceived benefit. The materials AI community has, in effect, traded epistemic flexibility for short-term predictive power, a trade that Jasanoff warns can become irreversible when technological systems and social orders co-evolve [2].

## Critique Point 4: The Illusion of Consensus

Critique point 4: The illusion of consensus arises when convergent narratives produced by model monoculture

create the appearance of widespread scientific agreement while systematically suppressing or rendering invisible genuine epistemic disagreement, uncertainty, and model sensitivity.

Messeri and Crockett have shown that AI systems in science generate a fluent but often illusory understanding [20], and materials AI provides a particularly clear case. Because the overwhelming majority of published results derive from closely related GNN variants trained on overlapping datasets, the literature appears to converge on robust findings—stable property predictions, reliable inverse-design candidates, and consistent rankings of material stability. Yet this apparent consensus is largely an artifact of the shared inductive biases and benchmark environments rather than independent corroboration. Disagreements that do exist are hidden in the supplementary materials or dismissed as hyperparameter sensitivity rather than fundamental model uncertainty.

Three materials-specific examples expose the illusion. First, in the prediction of formation energies for inorganic crystals, multiple GNN papers report near-identical error distributions on standard test sets, creating the narrative of “solved” accuracy. However, when the same architectures are evaluated on out-of-distribution disordered or high-pressure phases, performance collapses in ways that are rarely foregrounded. The convergent narrative, therefore, presents a false consensus on model reliability while the genuine disagreement—how much the graph representation actually captures thermodynamic reality—remains undiscussed.

Second, in the domain of thermal conductivity and phonon transport, GNN-based models are routinely benchmarked against one another using the same phonon datasets, yielding tightly clustered results that reviewers interpret as cross-validation. Yet Schütt *et al.*'s original SchNet work already demonstrated sensitivity to the choice of cutoff radius and basis functions [13], and later studies quietly confirm that small changes in graph construction can shift predicted transport properties by amounts larger than experimental uncertainty. The illusion of consensus is maintained by the convention of reporting only mean absolute errors within the dominant paradigm, thereby masking the fact that different research groups are effectively solving slightly different versions of the same problem.

Third, in generative discovery campaigns for novel perovskites or metal–organic frameworks, the community consensus appears to be that GNN-driven screening reliably identifies promising candidates. Yet when alternative metrics (synthesizability, long-term stability under operando conditions) are introduced, the ranked lists diverge dramatically. The seed critique already flagged this risk of illusory consensus [3], and subsequent scaling studies [25] continue to report headline numbers that reinforce the narrative without addressing the underlying model's sensitivity. Channing's analysis of AI for scientific discovery as a social problem is especially pertinent here [19]: the reward structure of publishing and funding favors consensus signals over transparent uncertainty quantification.

The illusion of consensus performs powerful gatekeeping work. It allows funding agencies to claim that “the community agrees” on prioritizing scaling GNNs over diversifying architectures. It permits reviewers to reject alternative approaches because “the field has already converged on superior methods.” And it discourages junior researchers from pursuing dissenting lines of inquiry because the apparent unanimity suggests that dissent is either uninformed or unproductive. The result is a scientific discourse that feels settled precisely because the mechanisms that would surface disagreement have been structurally disabled by the monoculture itself.

## Consequences for Materials Science

The convergence of scientific narratives around single-model regimes carries four interlocking consequences that threaten the long-term health of materials science as a discovery-oriented discipline.

### Consequence 1: Missed discoveries

Phenomena that fall outside the dominant graph-based representational grammar—such as emergent behavior in glassy systems, dynamic reconstruction at solid–liquid interfaces, or hierarchical self-assembly in bio-inspired composites—are never systematically explored because they are not legible to the dominant model. The epistemic narrowing and paradigm lock-in documented above ensure that entire regions of materials space remain unmapped.

## Consequence 2: Methodological stagnation

With research effort concentrated on incremental improvements within the GNN family, genuinely novel methodological paradigms receive neither funding nor talent. The review literature itself has narrowed [4, 5, 15, 23], creating a feedback loop in which the absence of alternative publications is taken as evidence that alternatives are unnecessary. Innovation in representation, rather than in scale or optimization, has effectively stalled.

## Consequence 3: Overconfidence

The illusion of consensus inflates trust in model outputs far beyond their actual robustness. When every high-profile paper reports similar success metrics within the same paradigm, practitioners and experimental collaborators begin to treat GNN predictions as near-oracles rather than as theory-laden approximations. This overconfidence risks experimental dead-ends and wasted resources when model-generated candidates fail in the laboratory for reasons the convergent narrative never anticipated.

## Consequence 4: Path dependency

Early architectural choices have now locked entire research trajectories for the next decade. Infrastructure built around GNNs—automated labs, shared databases, educational programs—will continue to channel future work along the same lines even if superior alternatives emerge. Reversing this path dependency will require coordinated intervention across funding, publishing, and training, precisely the kind of intervention that the current monoculture makes politically and cognitively difficult.

These consequences are not hypothetical; they are already visible in the slowed pace of truly unexpected materials discoveries despite the dramatic increase in computational throughput. The field risks becoming a highly efficient optimizer within a narrowly defined space rather than an explorer of the full materials universe.

## Alternative Approaches

To resist narrative convergence and restore epistemic diversity, five deliberate interventions are proposed.

**Table 2** contrasts the epistemic consequences of convergent narrative regimes with those of pluralistic

scientific systems in materials AI.

**Table 2.** Comparative epistemic effects of convergent narratives versus pluralistic scientific regimes in materials AI

Dimension	Convergent narrative regime (single-model)	Pluralistic regime (multi-model)	Epistemic outcome
Phenomenon coverage	Restricted to model-compatible systems	Expanded across heterogeneous systems	Increase in discovery potential
Methodological diversity	Low (intra-paradigm variation only)	High (orthogonal paradigms)	Robust inferences
Treatment of anomalies	Ignored or reframed	Investigated as signals	Paradigm evolution
Benchmark design	Bias-aligned	Adversarial and diverse	Reduced evaluation bias
Consensus formation	Apparent (artifact of similarity)	Contested and evidence-based	Genuine scientific agreement
Innovation trajectory	Incremental optimization	Conceptual exploration	Long-term advancement
Switching costs	High	Reduced via modularity	Increase in adaptability

### Alternative 1

**Deliberative Methodological Pluralism.** Funding agencies and journals should explicitly require and reward proposals and manuscripts that compare at least two genuinely orthogonal modeling paradigms on the same materials problem, with equal resources allocated to each. This shifts the burden of proof from the alternative to the dominant model and normalizes pluralism as a scientific virtue rather than an eccentricity.

### Alternative 2

**Adversarial Benchmarking.** Benchmark suites must include deliberately out-of-distribution test sets designed by researchers outside the dominant paradigm. They must report performance of multiple architectural families side-by-side without privileging any single baseline. Such adversarial construction would expose the hidden assumptions of current leaderboards and prevent the self-reinforcing cycle identified earlier.

### Alternative 3

**Narrative Diversity.** Journals and conference tracks should institute dedicated sections or sessions for “paradigm-challenging” work that explicitly questions the dominant story rather than extending it. Review criteria should reward clarity of dissent and methodological self-awareness rather than demanding immediate superiority over GNN baselines.

### Alternative 4: Paradigm auditing

Independent, rotating committees of materials scientists, philosophers of science, and AI ethicists should conduct periodic audits of the field’s modeling monoculture and publish public reports on citation concentration, benchmark bias, and epistemic blind spots. These audits would make the invisible costs of convergence visible and actionable.

### Alternative 5: Switching cost reduction

The community should invest in modular, interoperable toolkits that allow rapid migration between model families—standardized data schemas, universal featurizers, and automated translation layers—so that exploring an alternative architecture no longer requires rebuilding an entire experimental or computational pipeline. Reducing switching costs lowers the barrier to genuine exploration and weakens paradigm lock-in.

Taken together, these alternatives do not reject the undeniable power of data-driven methods; they seek to embed that power within a more reflexive, pluralistic scientific culture capable of self-correction.

## Conclusion

This critique has documented the emergence of model monoculture in materials AI, defined the problem of convergent scientific narratives, and elaborated four interlocking epistemic risks—epistemic narrowing,

suppression of alternatives, paradigm lock-in, and the illusion of consensus—that threaten the field’s long-term capacity for genuine discovery. The consequences are already material: missed opportunities in complex materials systems, methodological stagnation, and overconfidence in model outputs, and path-dependent trajectories that will be difficult to escape. Alternative approaches grounded in deliberative pluralism, adversarial benchmarking, narrative diversity, paradigm auditing, and deliberate reduction of switching costs offer practical pathways to resist convergence without sacrificing the benefits of modern machine learning.

The materials AI community stands at a Kuhnian inflection point where the choice between normal-science efficiency and revolutionary openness remains available—but only if the convergent narratives that currently dominate are actively contested. As Jasanoff reminds us, technological regimes are never inevitable; they are co-produced by technical choices and social practices. The self-referential warning embedded in the seed critique is therefore not merely diagnostic but prescriptive: the time to intervene is now, before single-model regimes harden into scientific orthodoxy. By embracing methodological and narrative diversity, materials science can retain its epistemic vitality while continuing to harness the predictive power of AI, ensuring that the next generation of materials discoveries is limited only by the creativity of the field rather than by the limitations of its dominant models.

## Acknowledgements

None

## Conflict of interest

None

## Financial support

None

## Ethics statement

None

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

- Suleimenov I, Gabrielyan O, Vitulyova Y. Dialectics of scientific revolutions from the point of view of innovations theory. *Wisdom*. 2022;4(24):25-35.
- Lee M. *The ethics of invention: Technology and the human future*. W W Norton Co.
- Petersen AM, Arroyave F, Pavlidis I. Methods for measuring social and conceptual dimensions of convergence science. *Res Eval*. 2023;32(2):256-72.
- Butler KT, Davies DW, Cartwright H, Isayev O, Walsh A. Machine learning for molecular and materials science. *Nature*. 2018;559(7715):547-55.
- 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.
- 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.
- Zunger A. Inverse design in search of materials with target functionalities. *Nat Rev Chem*. 2018;2(4):0121.
- Reiser P, Neubert M, Eberhard A, Torresi L, Zhou C, Shao C, et al. Graph neural networks for materials science and chemistry. *Commun Mater*. 2022;3(1):93.
- Zhong X, Gallagher B, Liu S, Kailkhura B, Hiszpanski A, Han TY. Explainable machine learning in materials science. *npj Comput Mater*. 2022;8(1):204.
- Liu Y, Yang Z, Zou X, Ma S, Liu D, Avdeev M, et al. Data quantity governance for machine learning in materials science. *Natl Sci Rev*. 2023;10(7):nwad125.
- Kleinberg J, Raghavan M. Algorithmic monoculture and social welfare. *Proc Natl Acad Sci*. 2021;118(22):e2018340118.
- 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.
- Schütt KT, Sauceda HE, Kindermans PJ, Tkatchenko A, Müller KR. SchNet—a deep learning architecture for molecules and materials. *J Chem Phys*. 2018;148(24).
- Gilmer J, Schoenholz SS, Riley PF, Vinyals O, Dahl GE. Neural message passing for quantum chemistry. In: *International Conference on Machine Learning*. PMLR; 2017. p. 1263-72.
- Schleder GR, Padilha AC, Acosta CM, Costa M, Fazzio A. From DFT to machine learning: Recent approaches to materials science—a review. *J Phys Mater*. 2019;2(3):032001.
- Wei J, Chu X, Sun XY, Xu K, Deng HX, Chen J, et al. Machine learning in materials science. *InfoMat*. 2019;1(3):338-58.
- Chong SS, Ng YS, Wang HQ, Zheng JC. Advances of machine learning in materials science: Ideas and techniques. *Front Phys*. 2024;19(1):13501.
- Jain A. Machine learning in materials research: Developments over the last decade and challenges for the future. *Curr Opin Solid State Mater Sci*. 2024;33:101189.
- Channing G, Ghosh A. AI for scientific discovery is a social problem. *Patterns*. 2026;7(3).
- Messeri L, Crockett MJ. Artificial intelligence and illusions of understanding in scientific research. *Nature*. 2024;627(8002):49-58.
- Ramprasad R, Batra R, Pilia G, Mannodi-Kanakkithodi A, Kim C. Machine learning in materials informatics: Recent applications and prospects. *npj Comput Mater*. 2017;3(1):54.

Oviedo F, Ferres JL, Buonassisi T, Butler KT. Interpretable and explainable machine learning for materials science and chemistry. *Acc Mater Res*. 2022;3(6):597-607.

Mobarak MH, Mimona MA, Islam MA, Hossain N, Zohura FT, Imtiaz I, et al. Scope of machine learning in materials research—A review. *Appl Surf Sci Adv*. 2023;18:100523.

Storey VC, Parsons J, Bueso AC, Tremblay MC, Lukyanenko R, Castillo A, et al. Domain knowledge in artificial intelligence: Using conceptual modeling to increase machine learning accuracy and explainability. *Data Knowl Eng*. 2025;160:102482.

Merchant A, Batzner S, Schoenholz SS, Aykol M, Cheon G, Cubuk ED. Scaling deep learning for materials discovery. *Nature*. 2023;624(7990):80-5.

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

Sauceda HE, Chmiela S, Poltavsky I, Müller KR, Tkatchenko A. Molecular force fields with gradient-domain machine learning: Construction and application to dynamics of small molecules with coupled cluster forces. *J Chem Phys*. 2019;150(11).

Korkmaz Y, Hegde D, Chintersingh KL, Alemohammad M, Kilic V, Flickinger MR, et al. Addressing data scarcity in materials science research with deep generative models. In: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*. IEEE; 2026. p. 1556-64.

Lake BM, Ullman TD, Tenenbaum JB, Gershman SJ. Building machines that learn and think like people. *Behav Brain Sci*. 2017;40:e253.