**ORIGINAL RESEARCH** **Open access**

# Validation Authority Erosion in High-Velocity Computational Materials Innovation

Ahmed El-Kholy[1*], Nour Abdelrahman[1], Karim Hassan[2]

## Abstract

The convergence of machine learning, high-throughput computation, and large-scale materials databases has propelled computational materials engineering into a regime of high-velocity innovation, where the generation of candidate structures and property predictions now occurs at rates orders of magnitude faster than traditional experimental validation. This shift has transformed the materials discovery pipeline from a sequential, experiment-centric process into a parallel, inference-dominated ecosystem. Yet the resulting disparity between computational throughput and empirical grounding has induced a subtle but profound erosion of validation authority—the epistemic weight traditionally assigned to direct experimental confirmation. This conceptual article synthesizes the computational and data-driven materials research landscape to examine how rapid inference challenges the established hierarchy of knowledge validation. Drawing on developments in machine learning interatomic potentials, uncertainty quantification, and autonomous discovery platforms, the analysis reveals systemic pressures that redistribute authority across data, models, and discovery outputs. To address these dynamics, the Velocity-Induced Validation Authority Reconfiguration (VIVAR) Framework is introduced as an original systems-level architecture. VIVAR conceptualizes validation not as a static endpoint but as a dynamic, reconfigurable layer embedded within the discovery pipeline. It delineates structural layers, forward-propagating data-to-discovery flows, bidirectional feedback mechanisms, and computational steering logics that enable adaptive authority allocation. By interpreting validation authority as an infrastructure resource subject to erosion and realignment, the framework provides interpretive tools for managing epistemic risk and infrastructure trade-offs in accelerated materials ecosystems. The implications extend beyond individual workflows to the broader architecture of computational materials innovation, offering a lens for designing platforms that sustain discovery velocity while preserving epistemic integrity. In an era where computational predictions increasingly precede and sometimes supplant experimentation, such reconfiguration becomes essential for the sustainable advancement of the field.

**Keywords** Uncertainty quantification, Discovery pipelines, Data-driven materials discovery, Validation authority, High-velocity computation, Epistemic reconfiguration

*Correspondence:
Ahmed El-Kholy
ahmed.elkholy@gmail.com
[1] Department of Computational Materials Engineering, Faculty of Engineering, Alexandria University, Alexandria, Egypt
[2] Department of Materials Data Analytics, Faculty of Engineering, Ain Shams University, Cairo, Egypt

# Introduction

## The acceleration of computational materials innovation

Computational materials science has undergone a profound transformation since the mid-2010s, driven by the integration of machine learning, large-scale databases, and automated workflows. Early frameworks for property prediction using compositional descriptors evolved into sophisticated graph neural networks and foundation models

capable of screening millions of candidates in days rather than years [1-4]. Initiatives such as the Materials Project have scaled density functional theory calculations across vast chemical spaces, while machine learning interatomic potentials now enable molecular dynamics simulations at near-DFT accuracy and orders-of-magnitude greater efficiency [5, 6]. These advances have shifted the bottleneck in materials innovation from generation of hypotheses to their validation and selection.

High-velocity innovation manifests in multiple dimensions: the sheer volume of predictions generated per unit time, the diversity of modalities (electronic, mechanical, thermal, optical) addressed simultaneously, and the integration of generative models that propose entirely new compositions and structures [7, 8]. Publications in leading journals document workflows that traverse from elemental composition to predicted device performance within unified computational pipelines [9, 10]. This velocity is not merely quantitative; it alters the temporal rhythm of discovery, compressing what were once multi-year experimental campaigns into iterative computational cycles measured in weeks or even hours.

## The traditional role of experimental validation

Historically, experimental validation functioned as the ultimate epistemic authority in materials research. Synthesis, characterization, and performance testing provided the ground truth against which theoretical models were judged [11]. This hierarchy ensured that computational tools served primarily as hypothesis generators or interpretive aids, with authority flowing unidirectionally from experiment to model refinement. Uncertainty in computational outputs was resolved through direct comparison with measured properties, reinforcing the primacy of empirical data [12, 13].

Even as computational methods matured, this structure persisted. Density functional theory calculations required experimental benchmarks for functional selection, and early machine learning models were trained and validated predominantly on experimental datasets [14, 15]. The literature reflects a consistent emphasis on closing the loop between prediction and measurement, with validation serving as both a quality gate and a source of new training data [16, 17].

## Emerging tensions in the data-driven paradigm

The scaling of computational capacity has introduced structural tensions. Machine learning potentials and foundation models now produce predictions whose volume exceeds feasible experimental follow-up by several orders of magnitude [4]. Uncertainty quantification techniques, while essential, often reveal confidence intervals that are themselves model-dependent and difficult to anchor without extensive experimentation [18-20]. High-throughput screening campaigns generate thousands of promising candidates, yet only a tiny fraction can be synthesized and tested within realistic resource constraints [9, 21].

These pressures manifest as a gradual erosion of validation authority. Computational outputs increasingly circulate in the literature and databases with provisional acceptance, sometimes supported by internal consistency checks or cross-model agreement rather than direct experimental confirmation [22, 23]. The epistemic weight once concentrated in experimental results is diffusing across computational ensembles, surrogate models, and literature-derived embeddings [24]. This redistribution is not necessarily detrimental but represents a reconfiguration whose dynamics remain underexamined at the systems level.

The literature synthesis reveals recurring motifs: the need for robust uncertainty quantification to compensate for reduced experimental sampling [18, 19], the emergence of autonomous laboratories that blend computation and robotics yet still confront validation bottlenecks [23], and calls for new standards in data-driven materials science that acknowledge the changing role of validation [22]. Yet these contributions remain fragmented, focused on specific techniques or applications rather than the overarching epistemic infrastructure.

## Positioning the VIVAR framework

Against this backdrop, the present work advances the Velocity-Induced Validation Authority Reconfiguration (VIVAR) Framework as a conceptual systems architecture tailored to high-velocity computational materials innovation. VIVAR interprets validation authority as a configurable resource within the discovery pipeline, subject to erosion under velocity pressures but amenable to deliberate realignment through structured feedback and steering logics. The framework integrates insights from the

computational materials literature into a unified interpretive model that emphasizes pipeline dynamics, representation-inference interactions, and epistemic risk structures. It positions validation not as a fixed gate but as an adaptive layer that can be dynamically allocated across computational and experimental modalities to sustain both velocity and integrity.

# Theoretical Background & Literature Synthesis

## Machine learning-driven property prediction and its scaling

The foundational shift toward data-driven materials engineering rests on the development of machine learning architectures capable of learning structure-property relationships from large datasets. Early approaches using crystal graph convolutional networks and elemental composition models demonstrated that accurate predictions could be achieved even with limited experimental data when augmented by high-fidelity computations [2, 10]. Subsequent scaling to foundation models has further accelerated this capability, enabling screening across millions of candidates with remarkable transferability [4].

These methods have fundamentally altered the materials design space. Rather than relying on intuition-driven selection of a few dozen compositions, researchers now explore combinatorial libraries that were previously intractable [7, 9]. The literature documents successful applications in perovskites, thermoelectrics, and battery materials, where computational predictions have guided experimental efforts with increasing precision [21, 25]. Yet this success has also highlighted limits: models trained primarily on computational data risk propagating systematic errors, and the velocity of prediction often outstrips the capacity to generate diverse, high-quality training examples [14, 15].

## Uncertainty quantification as a compensatory mechanism

Recognition of these limits has driven substantial investment in uncertainty quantification (UQ) for computational materials models. Dropout-based neural network potentials and ensemble methods now provide principled estimates of prediction confidence, allowing researchers to flag regions of high epistemic uncertainty for prioritized experimental attention [18, 19]. Multi-fidelity approaches further refine this by combining low-cost screening with targeted high-accuracy calculations [12].

UQ serves a dual role in the high-velocity regime. It acts as an internal validation proxy when external experimentation is constrained and as a steering signal for directing resources toward the most impactful measurements [16]. The literature shows growing sophistication in UQ, from Bayesian methods in active learning to graph-network-based confidence metrics [3, 26]. Nevertheless, UQ remains model-bound; its reliability depends on the representativeness of the training distribution and cannot fully substitute for direct empirical grounding in novel chemical spaces [20].

## High-throughput ecosystems and autonomous discovery

Parallel to advances in individual models, the field has developed integrated ecosystems that automate the entire discovery pipeline. The Materials Project exemplifies infrastructure-scale efforts that combine databases, workflows, and community contributions to sustain continuous computational exploration [5]. Autonomous platforms incorporating robotic synthesis and characterization further close the loop, yet even these systems reveal persistent validation challenges when confronted with truly novel materials [23].

Literature on these ecosystems emphasizes the importance of standardized data practices and interoperability [22]. They also reveal a recurring pattern: the faster the computational front end, the greater the pressure on downstream validation stages. Active learning strategies attempt to mitigate this by adaptively selecting experiments that maximize information gain, yet the overall throughput mismatch remains a structural feature of the paradigm [13, 16].

## Representation–inference interactions and epistemic risk

A deeper theme emerging from the synthesis concerns how different representations (compositional, structural, graph-based) interact with inference engines to shape epistemic outcomes. Graph networks and transformer architectures excel at capturing local and global patterns but can produce confident yet incorrect extrapolations in under-sampled

regions [3, 4]. Literature-derived embeddings, such as those from unsupervised analysis of scientific text, add another layer of knowledge but introduce risks of inherited biases or incomplete coverage [24].

These interactions generate epistemic risk structures that are distinct from traditional modeling uncertainties. In high-velocity environments, the rapid cycling of predictions can amplify small representational biases into widespread acceptance of provisional results. The literature calls for new approaches to traceability and provenance that track how authority accrues across the pipeline [22, 27].

## Validation in the computational era: Toward reconfiguration

Collectively, the body of work documents a field in transition. Computational methods have achieved remarkable velocity and scope, yet the mechanisms for maintaining epistemic integrity have not scaled commensurately. Validation authority is eroding not through deliberate devaluation of experiment but through the sheer momentum of computational output. The VIVAR Framework builds on these insights to propose a systems-level response that treats this erosion as a manageable dynamic rather than an inevitable loss.

## Proposed conceptual framework

### The Velocity-Induced Validation Authority Reconfiguration (VIVAR) framework

The VIVAR Framework conceptualizes materials discovery as a multi-layered epistemic pipeline in which validation authority is dynamically reconfigured in response to computational velocity. It comprises four structural layers connected by forward data-to-discovery pipelines, bidirectional feedback loops, and explicit steering logics. Authority is treated as a distributed resource that can be allocated, eroded, or reinforced depending on the velocity-pressure profile of the workflow.

The layers are:

·**Data Reservoir Layer**: Aggregates heterogeneous inputs from computational databases, literature embeddings, and sparse experimental records. This layer emphasizes curation and provenance tracking to maintain traceability.

·**Inference Engine Layer**: Hosts machine learning models, interatomic potentials, and generative architectures that produce property predictions and candidate structures at high velocity.

·**Authority Nexus Layer**: The core reconfiguration zone where validation authority is assessed and allocated. This layer integrates uncertainty quantification outputs with pipeline context to decide the balance between computational confidence and the need for experimental grounding.

·**Discovery Output Layer**: Translates validated predictions into actionable material candidates, application concepts, or new hypotheses, while generating feedback signals for upstream layers.

**Data → Model → Discovery Pipelines**

Forward pipelines propagate information unidirectionally from the Data Reservoir through the Inference Engine to the Discovery Output. At each transition, authority is provisionally assigned based on internal consistency metrics. In low-velocity regimes, the pipeline may route outputs through the Authority Nexus for experimental invocation. In high-velocity operation, the pipeline accelerates by weighting computational authority more heavily when uncertainty falls below configurable thresholds.

### Feedback loops

Bidirectional loops enable continuous recalibration. Post-discovery feedback from the Output Layer returns to the Authority Nexus to adjust allocation weights and to the Data Reservoir to enrich training distributions. A secondary loop connects the Inference Engine directly to the Authority Nexus, allowing model ensembles to self-assess and request targeted external validation when internal disagreement exceeds defined bounds.

### Computational steering logics

Steering logics operate at the Authority Nexus as decision functions that modulate pipeline velocity and authority distribution. These logics incorporate both local uncertainty signals and global pipeline state, enabling adaptive throttling of computation when validation risk accumulates or acceleration when confidence is high.

The dynamics of authority reconfiguration within VIVAR can be conceptualized through three interrelated expressions.

First, the erosion of validation authority under velocity pressure may be expressed as:

$$\frac{dA}{dt} = -\kappa \left( Vc - \lambda Et \right) \quad (1)$$

where A represents the instantaneous validation authority index, $V_c$ denotes computational velocity (predictions per unit time), $E_t$ is experimental throughput capacity, κ is a system-specific erosion coefficient reflecting infrastructure rigidity, and λ is a latency scaling factor. This form captures the interpretive insight that authority erodes when computational output velocity persistently exceeds the rate at which empirical grounding can be supplied, with the magnitude modulated by the latency between prediction and possible validation.

Second, the feedback-mediated realignment of authority across the pipeline can be conceptualized as:

$$A_{t+1} = At + \mu \left( Df \cdot Ur - 1 \right) \quad (2)$$

Here, $A_{t+1}$ is the updated authority distribution, $D_f$ represents the information gain from discovery-stage feedback, $U_r$ denotes residual unresolved uncertainty, and μ \mu μ is the reconfiguration responsiveness of the system. This expression illustrates how feedback loops counteract erosion by amplifying authority adjustments in regions of high uncertainty reduction, thereby stabilizing the epistemic pipeline without sacrificing velocity.

Third, the steering logic for modulating pipeline flow at the Authority Nexus may be represented as:

$$S = argmaxw \left( Pd \otimes Wa \right) \quad (3)$$

where S is the steering decision vector (e.g., accelerate, invoke experiment, or recalibrate model), $P_d$ is the vector of prediction confidences across candidates, $W_a$ is the current authority weight matrix, and ⊗ denotes a weighted interaction operator. This formulation interprets steering as an optimization over the interaction between confidence and authority allocation, ensuring that velocity is preserved while epistemic risk remains bounded. As conceptualized in **Figure 1**, validation authority is dynamically reconfigured across layered discovery infrastructures through feedback-mediated steering logics operating under computational velocity pressure.

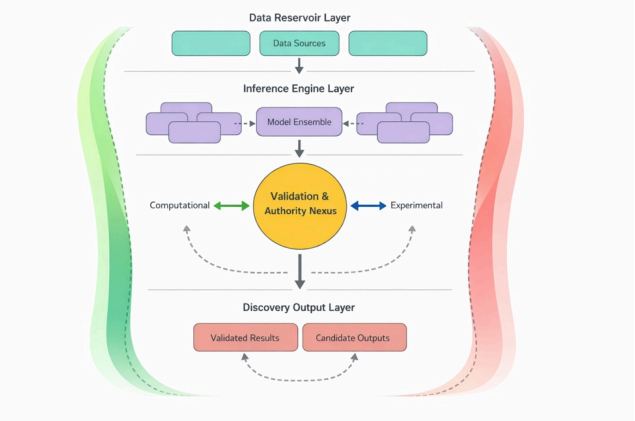


**Figure 1.** Velocity-Induced validation authority reconfiguration architecture.

A multi-layered systems diagram illustrating how validation authority is redistributed across high-velocity computational materials discovery pipelines. The architecture comprises four structural layers—Data Reservoir, Inference Engine, Authority Nexus, and Discovery Output—connected by forward prediction pipelines, bidirectional feedback loops, and adaptive steering logics. The Authority Nexus operates as the central reconfiguration node, dynamically allocating epistemic weight between computational inference and experimental grounding based on uncertainty signals and velocity pressure. Dashed propagation pathways depict feedback-mediated authority realignment, while steering vectors illustrate adaptive routing decisions including computational acceleration, experimental invocation, and model recalibration. The diagram conceptualizes validation not as a terminal gate but as an embedded, dynamically modulated infrastructure layer. The layered epistemic architecture of VIVAR is structurally decomposed in **Table 1**.

**Table 1.** Structural Layers and Functional Roles in the VIVAR Framework

| Framework Layer | Primary Function | Epistemic Role | Authority Dynamics |
|---|---|---|---|
| Data Reservoir | Aggregates heterogeneous inputs | Knowledge provenance anchor | Authority seeding |

| | | | |
|---|---|---|---|
| Inference Engine | Generates predictions & candidates | Computational inference core | Authority amplificatio |
| Authority Nexus | Allocates validation weight | Epistemic governance center | Authority redistributio |
| Discovery Output | Produces validated materials insights | Translational realization layer | Authority manifestatic |

## Analytical implications

The VIVAR Framework reframes several core operational challenges in computational materials engineering as matters of infrastructure configuration rather than isolated methodological shortcomings. At the systems level, it implies that high-velocity pipelines must embed authority reconfiguration as a native capability rather than an add-on validation stage. This has direct consequences for workflow architecture: the Authority Nexus becomes the primary control point where velocity pressure is converted into adaptive decisions, replacing rigid sequential gates with context-aware routing [16, 22, 23].

A central implication concerns the management of epistemic risk structures. In conventional pipelines, risk accumulates downstream; in VIVAR, risk is actively redistributed upstream through feedback loops that continuously enrich the Data Reservoir with authority-weighted observations. This redistribution prevents the formation of authority debt—accumulated provisional predictions that later require costly retroactive grounding. The framework thus interprets data curation not as passive archiving but as an active epistemic maintenance process [4, 5, 24].

Representation–inference interactions receive particular emphasis. Different model architectures (graph networks, transformers, interatomic potentials) generate distinct authority signatures. The framework implies that pipeline designers should treat these signatures as configurable parameters within the Authority Nexus, enabling hybrid ensembles where authority is allocated proportionally to representational strengths in specific chemical subspaces [2, 3, 8].

To capture the trade-off dynamics between computational velocity and sustained epistemic integrity, the interaction can be conceptualized as:

$$\begin{matrix}\eta \\ = Vc \\ \cdot An1 \\ + \beta\left(\frac{Vc}{\mathrm{Emax}-1}\right)2\end{matrix} \quad (4)$$

where η denotes effective discovery efficiency, $V_c$ is instantaneous computational velocity, $A_n$ is the normalized authority level at the Nexus, Emax is the maximum experimental anchoring capacity of the ecosystem, and β \beta β is a damping coefficient reflecting infrastructure inertia. This expression illustrates how efficiency peaks when velocity and authority are balanced, decaying quadratically when velocity exceeds the system's capacity to maintain authority through feedback and steering.

Another implication is the emergence of new roles for uncertainty quantification. Rather than serving solely as a confidence metric, UQ outputs become inputs to steering logics that modulate the entire pipeline [18-20]. This elevates UQ from a diagnostic tool to a core governance mechanism within the discovery infrastructure.

Finally, the framework highlights infrastructure trade-offs at the community level. Adoption of VIVAR-like architectures would require shared standards for authority provenance metadata, enabling interoperability across independent platforms and preventing fragmentation of epistemic resources [9, 22]. These implications collectively position validation authority as a design variable rather than a fixed constraint, offering a pathway to scale discovery velocity without proportional increases in experimental infrastructure.

# Results and Discussion

The VIVAR Framework integrates and extends the conceptual threads running through computational materials literature. It synthesizes the scaling successes of machine learning potentials and foundation models [4, 6] with the persistent validation bottlenecks documented in high-throughput campaigns [9, 21] and the growing emphasis on uncertainty-aware workflows [18, 19, 26].

Where prior work has focused on improving individual components—better models, larger datasets, more sophisticated UQ—the framework shifts attention to the connective tissue: the reconfiguration mechanisms that allow these components to operate coherently under velocity pressure.

This systems-level perspective reveals that authority erosion is not an accidental byproduct but a structural feature of the current paradigm. The literature consistently shows that computational throughput now exceeds experimental capacity by factors of $10^2$–$10^4$ [4, 5], yet most studies still treat validation as an external, rate-limiting step. VIVAR internalizes this mismatch, transforming it into a controllable dynamic through layered feedback and steering. In doing so, it aligns with emerging calls for new standards in data-driven materials science [22] while providing an operational architecture to realize them.

Implementation challenges remain interpretive rather than prescriptive. The Authority Nexus, while conceptually modular, would require careful calibration of steering logics to avoid oscillatory behavior in which the system over-corrects between acceleration and experimental invocation. Similarly, the feedback loops depend on robust provenance tracking; without it, authority signals could propagate errors rather than corrections [22, 24]. These challenges underscore that VIVAR is not a plug-and-play solution but a conceptual scaffold for platform developers and standards bodies.

Compared with existing frameworks, VIVAR differs in its explicit treatment of validation authority as a distributed, time-varying resource. Active learning approaches [16] optimize experiment selection but do not address the broader pipeline authority dynamics. Autonomous laboratory concepts [23] close the loop at the hardware level but rarely model the epistemic reconfiguration required at scale. The framework thus fills a conceptual gap at the infrastructure layer, where most innovation pressure is currently concentrated.

Broader implications extend to the sociology of the field. As computational predictions increasingly circulate with provisional authority, the community may need to develop new norms for citing and building upon results whose validation status is explicitly quantified and dynamically updated. This could accelerate discovery in sparse-data regimes—such as complex oxides or high-entropy alloys—where traditional validation is fundamentally rate-limited [7, 21]. At the same time, it raises the importance of maintaining human oversight within the steering logics to preserve scientific judgment in high-stakes applications.

The framework also suggests pathways for cross-domain transfer. Similar velocity-authority tensions appear in adjacent fields such as autonomous chemical synthesis and generative biology. The layered structure of VIVAR, with its emphasis on reconfigurable authority, offers a transferable template for epistemic infrastructure design beyond materials science.

Ultimately, the discussion centers on sustainability. High-velocity innovation is already here; the question is whether the supporting epistemic infrastructure can evolve quickly enough to sustain it. The VIVAR Framework provides one interpretive model for ensuring that it does.

# Conclusion

The Velocity-Induced Validation Authority Reconfiguration (VIVAR) Framework conceptualizes the ongoing transformation in computational materials engineering as a manageable reconfiguration of epistemic infrastructure. By positioning validation authority as a dynamic layer within the discovery pipeline—subject to erosion under velocity pressure yet restorable through structured feedback and steering—the framework offers a systems-level response to the central tension of the data-driven era.

The analysis demonstrates that the disparity between computational generation and empirical grounding is not a temporary limitation but a defining characteristic of modern materials innovation. The proposed architecture addresses this by redistributing authority across data, models, and outputs in a manner that preserves both velocity and integrity. Its layered structure, forward pipelines, bidirectional loops, and steering logics provide a coherent interpretive model for designing platforms that operate effectively at the frontier of chemical and structural complexity.

As the field continues to scale toward foundation models, autonomous ecosystems, and real-time discovery, frameworks that explicitly model epistemic dynamics will become essential infrastructure components. VIVAR contributes to this evolution by shifting the discourse from isolated methodological advances to the connective logics that determine whether high-velocity computation translates into reliable materials progress.

The ultimate value of such reconfiguration lies in its capacity to sustain the transformative promise of computational materials science while safeguarding the epistemic standards that define scientific knowledge.

## Acknowledgements

None

## Conflict of interest

None

## Financial support

None

## Ethics statement

None

Received: 11 Jan 2025 Revised: 25 Feb 2025 Accepted: 21 May 2025
Published online: 18 September 2025



## References

Ramprasad R, Batra R, Pilania G, Mannodi-Kanakkithodi A, Kim C. Machine learning in materials informatics: Recent applications and prospects. npj Comput Mater. 2017;3(1):54. https://doi.org/10.1038/s41524-017-0056-5.

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. https://doi.org/10.1103/PhysRevLett.120.145301.

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. https://doi.org/10.1021/acs.chemmater.9b01294.

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.

Horton MK, Huck P, Yang RX, Munro JM, Dwaraknath S, Ganose AM, et al. Accelerated data-driven materials science with the materials project. Nat Mater. 2025;24(10):1522-32. https://doi.org/10.1038/s41563-025-02272-0.

Mishin Y. Machine-learning interatomic potentials for materials science. Acta Mater. 2021;214:116980. https://doi.org/10.1016/j.actamat.2021.116980.

Tao Q, Xu P, Li M, Lu W. Machine learning for perovskite materials design and discovery. npj Comput Mater. 2021;7(1):23.

Rosen AS, Sahasrabuddhe S, Jain A. High-throughput prediction of material properties using compositionally-restricted attention-based network. npj Comput Mater. 2022;8(1):1-11.

Choudhary K, DeCost B, Tavazza F. Machine learning for high-throughput materials discovery and design. Comput Mater Sci. 2022;210:111437.

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. https://doi.org/10.1038/s41598-018-35934-y.

Zunger A. Inverse design in search of materials with target functionalities. Nat Rev Mater. 2018;3(6):1-15.

Pilania G, Gubernatis JE, Lookman T. Multi-fidelity machine learning models for accurate bandgap predictions of solids. Comput Mater Sci. 2017;129:156-63.

Li W, Jacobs R, Morgan D. Predicting the thermodynamic stability of solids combining density functional theory and machine learning. Comput Mater Sci. 2018;150:454-63.

Schmidt J, Marques MRG, Botti S, Marques MAL. Recent advances and applications of machine learning in solid-state materials science. npj Comput Mater. 2019;5(1):83.

Himanen L, Geurts A, Foster AS, Rinke P. Data-driven materials science: status, challenges, and perspectives. Adv Sci. 2019;6(21):1900808. https://doi.org/10.1002/advs.201900808.

Lookman T, Balachandran PV, Xue D, Yuan R. Active learning in materials science with emphasis on adaptive sampling using uncertainties for targeted design. npj Comput Mater. 2019;5(1):21. https://doi.org/10.1038/s41524-019-0153-8.

Oviedo F, Ren Z, Sun S, Settens C, Liu Z, Liu Y, et al. Fast and interpretable classification of small X-ray diffraction datasets using data augmentation and deep neural networks. npj Comput Mater. 2019;5(1):60. https://doi.org/10.1038/s41524-019-0196-x.

Wen M, Tadmor EB. Uncertainty quantification in molecular simulations with dropout neural network potentials. npj Comput Mater. 2020;6(1):124. https://doi.org/10.1038/s41524-020-00390-8.

Bilbrey JA, Firoz JS, Lee MS, Choudhury S. Uncertainty quantification for neural network potential foundation models. npj Comput Mater. 2025;11:109. https://doi.org/10.1038/s41524-025-01572-y.

Swinburne TD, Marinica MC. Uncertainties in the prediction of defect properties using machine-learned interatomic potentials. npj Comput Mater. 2024;10:45.

Wang Z, Zhang Z, Liu J, Wang Z, Li H, Wang Y, et al. Machine learning-accelerated discovery of high-performance thermoelectric materials. Adv Mater. 2023;35(12):2209123.

Butler KT, Choudhary K, Csanyi G, Ganose AM, Kalinin SV, Morgan D. Setting standards for data driven materials science. npj Comput Mater. 2024;10:231. https://doi.org/10.1038/s41524-024-01411-6.

Kalinin SV, Ziatdinov M, Hinkle J, Jesse S, Ghosh A, Kelley KP, et al. Automated and autonomous experiments in electron and scanning probe microscopy. npj Comput Mater. 2021;7(1):1-15.

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. https://doi.org/10.1038/s41586-019-1335-8.

Wang H, Ji Y, Li Y. Machine learning for materials design and discovery. Adv Mater. 2022;34(36):2104113.

Dunn A, Wang Q, Ganose A, Dopp D, Jain A. Benchmarking materials property prediction methods: The Matbench test set and Automatminer reference algorithm. npj Comput Mater. 2020;6(1):138.

Morgan D, Jacobs R. Opportunities and challenges for machine learning in materials science. Annu Rev Mater Res. 2020;50:71-103. https://doi.org/10.1146/annurev-matsci-070218-010015.