

ORIGINAL RESEARCH

Open access

Temporal Governance Gaps: Oversight Latency in Accelerated Materials Discovery Systems

Nguyen Thanh Huy^{1*}, Pham Quang Minh¹, Le Thi Bich²

Abstract

The integration of computational modelling, machine learning, and robotic automation has fundamentally altered the tempo of materials discovery. High-throughput density functional theory databases, graph neural networks trained on vast materials corpora, and self-driving laboratories now generate and evaluate candidate structures at rates orders of magnitude beyond conventional workflows. These systems excel at navigating combinatorial spaces and proposing materials with targeted properties, yet the very acceleration they enable exposes a structural vulnerability: oversight latency. Oversight here denotes the epistemic processes—validation against physical reality, uncertainty propagation, causal interpretation, and knowledge consolidation—that anchor computational predictions within reliable materials engineering practice. When discovery pipelines advance faster than these processes can respond, temporal governance gaps emerge. Unvalidated or partially validated candidates propagate through downstream design, risking cascading epistemic errors in applications ranging from energy storage to quantum materials. This article synthesizes the literature on accelerated platforms articulate oversight latency as a systemic, rather than incidental, feature of contemporary data-driven ecosystems. We introduce the Temporal Governance Synchronization Framework (TGSF), an original conceptual architecture that reframes discovery pipelines as coupled dynamical systems whose synchronization determines epistemic integrity. TGSF identifies structural layers, feedback topologies, and steering logics that can align discovery velocity with governance capacity without sacrificing throughput. By foregrounding temporal dynamics, the framework offers infrastructure-level guidance for designing next-generation materials acceleration platforms that are both rapid and epistemically robust. Its implications extend to the sustainable scaling of computational materials engineering and the responsible stewardship of autonomous discovery systems.

Keywords Self-driving laboratories, Data-driven pipelines, Temporal governance, Oversight latency, Accelerated materials discovery, Epistemic synchronization

*Correspondence:

Nguyen Thanh Huy
huy.nguyen@gmail.com

¹ Department of Materials Data Science, Faculty of Engineering, Vietnam National University, Hanoi, Vietnam

² Department of Computational Engineering Systems, Faculty of Engineering, Can Tho University, Can Tho, Vietnam

Introduction

The acceleration imperative in materials engineering

Computational materials science has undergone a structural transformation, evolving from an auxiliary analytical instrument into the primary engine of materials discovery. Historically, computation served to rationalize experimentally observed phenomena—predicting phase stability, electronic structure, or defect energetics in support of laboratory synthesis. Over the past decade, however, this relationship has inverted. Discovery pipelines are now initiated computationally, with simulation and machine learning directing experimental validation rather than merely interpreting it.

High-throughput virtual screening exemplifies this inversion. Early combinatorial simulations explored thousands of candidate compositions; contemporary infrastructures routinely evaluate millions through automated density functional theory (DFT) workflows augmented by surrogate modeling architectures [1-3]. Distributed computing frameworks, automated convergence protocols, and error-correcting workflows have rendered large-scale first-principles evaluation both reproducible and scalable. Parallel advances in machine learning have further amplified throughput. Graph neural networks, kernel methods, and deep representation learners trained on curated repositories now approximate DFT-level accuracy at orders-of-magnitude lower computational cost [4, 5].

Large open repositories—including the Materials Project and related federated databases—serve as epistemic backbones for this transformation. Their standardized datasets, interoperable APIs, and continuously expanding property annotations provide training corpora that enable rapid model generalization across compositional and structural spaces. The result is an infrastructure in which predictive inference operates at near-simulation fidelity while retaining the speed of statistical estimation.

This quantitative leap in evaluative capacity carries qualitative implications. Traditional discovery cycles unfolded across human timescales: hypothesis generation, synthesis, and validation progressed sequentially over months or years. In contrast, contemporary platforms generate candidate materials, simulate properties, and propose synthesis pathways within hours or days [6-8]. Materials engineering has thus transitioned from a regime of data scarcity to one of data deluge, from serial experimentation to massively parallel exploration, and from human-paced reasoning to machine-paced hypothesis production.

From high-throughput computation to autonomous discovery

A second wave of acceleration has emerged through the convergence of computational prediction with robotic execution. Closed-loop automation infrastructures—commonly termed self-driving laboratories—integrate robotic synthesis platforms, in-situ characterization instrumentation, and active learning algorithms into unified experimental ecosystems [6, 7, 9-12]. Within these systems, machine learning models propose candidate materials; robotic platforms synthesize them; real-time diagnostics characterize outcomes; and resulting data are reintegrated into model retraining cycles.

Empirical demonstrations illustrate the maturity of this paradigm. Autonomous thin-film deposition platforms and inorganic synthesis systems reported by MacLeod *et al.* [6] and Szymanski *et al.* [7] achieve end-to-end experimental iteration without continuous human oversight. At the computational frontier, Merchant *et al.* [1] demonstrate that large-scale deep learning architectures can propose thermodynamically stable inorganic compounds across compositional spaces previously inaccessible to systematic exploration.

Efficiency gains are profound. Discovery metrics—candidate evaluations per unit time, experimental cycles closed per day, optimization convergence rates—have become central performance indicators. Autonomous platforms compress design–make–test–analyze loops into tightly coupled computational–experimental feedback systems, dramatically reducing the latency between prediction and empirical observation.

Yet prevailing discourse remains throughput-centric. System performance is predominantly measured through volumetric productivity rather than epistemic robustness. Less analytical attention has been directed toward the temporal structure of discovery itself—specifically, the relative velocities at which prediction, validation, interpretation, and knowledge consolidation occur.

The unexamined dimension: Temporal aspects of governance

Governance within accelerated discovery infrastructures is not solely regulatory or institutional; it is intrinsically epistemic. It encompasses the procedural and infrastructural mechanisms through which computational communities ensure that proposed materials are physically plausible, synthetically accessible, mechanistically interpretable, and reproducible within established structure–property paradigms. Governance operates through uncertainty quantification protocols,

experimental cross-validation regimes, causal attribution analyses, and the integration of validated findings into shared scientific knowledge systems.

While literature on autonomous discovery acknowledges validation bottlenecks and reproducibility constraints [8, 9, 13, 14], these discussions are often framed as logistical or technical limitations rather than systemic temporal misalignments. Experimental validation pipelines frequently lag behind computational or robotic proposal engines by weeks to months. Characterization throughput, synthesis feasibility constraints, and interpretive analysis introduce unavoidable latency into the adjudication of computational claims.

In high-velocity discovery regimes, such latency produces operational exposure windows. During these intervals, unvalidated or partially validated candidates may influence subsequent model training, active learning prioritization, or downstream design decisions. Closed-loop optimization systems, in particular, risk recursively amplifying preliminary or uncertain signals.

This condition constitutes a temporal governance gap: a structural misalignment between the rate at which discovery systems generate knowledge claims and the rate at which those claims can be responsibly verified, contextualized, and epistemically consolidated.

Epistemic consequences of governance latency

Temporal governance gaps are not merely operational inefficiencies; they reshape the epistemic topology of the field. Accelerated iteration can amplify latent biases embedded within training corpora [5], particularly where compositional or structural coverage is uneven. Surrogate models may propagate uncertainty fields inherited from sparse simulation data or approximated descriptors [15, 16]. Active learning policies, optimized for rapid convergence, may privilege synthesizable yet suboptimal candidates over theoretically superior but experimentally demanding alternatives.

Over time, these dynamics can generate self-reinforcing discovery trajectories. Computational consensus may crystallize around rapidly validated material classes, while slower-to-validate domains remain underexplored. In extreme scenarios, the field risks entering a state of epistemic momentum—a condition in which the velocity of computational agreement outpaces the accumulation of empirical grounding. Under such conditions, discovery pipelines may optimize within epistemically constrained subspaces while maintaining the illusion of global exploration.

Positioning a synchronization perspective

Existing methodological frameworks address elements of this challenge. Optimization architectures accelerate search efficiency; uncertainty-aware exploration strategies attempt to balance risk and novelty; multi-objective design frameworks integrate performance, stability, and manufacturability constraints [13, 15, 17–20]. These contributions are foundational to autonomous discovery science.

However, they conceptualize latency primarily as a parameter to be minimized—an engineering inefficiency to be compressed through faster computation, improved robotics, or denser datasets. Such a framing overlooks latency's structural role within epistemic governance. Validation delay, interpretive deliberation, and knowledge assimilation are not reducible frictions but constitutive features of responsible discovery.

The Temporal Governance Synchronization Framework (TGSF) reframes accelerated materials innovation as a system of coupled temporal pipelines. Rather than optimizing velocity in isolation, TGSF models discovery infrastructures through the alignment—or misalignment—of data generation, predictive inference, experimental validation, and epistemic consolidation cycles. By analyzing latency distributions, feedback topologies, and steering logics, the framework identifies synchronization thresholds necessary to preserve epistemic integrity under conditions of extreme acceleration.

In doing so, TGSF provides a conceptual scaffold for designing next-generation discovery infrastructures—systems capable not only of unprecedented speed but of temporally harmonized governance. The sections that follow develop the theoretical foundations of this perspective, synthesize relevant literature on accelerated materials ecosystems, and articulate the synchronization framework in structural detail.

Theoretical Background & Literature Synthesis

Computational foundations of accelerated screening

Early acceleration in computational materials discovery relied on systematic enumeration and high-throughput quantum calculations. The development of large-scale repositories such as the Open Quantum Materials Database enabled thermodynamic stability screening across vast compositional spaces [3]. These infrastructures demonstrated the feasibility of evaluating millions of hypothetical compounds but simultaneously exposed a persistent translational bottleneck: the conversion of computational predictions into experimentally validated materials.

Subsequent methodological evolution integrated machine learning to prioritize candidate selection. Foundational reviews demonstrated that supervised learning models trained on density functional theory datasets could predict formation energies, electronic band gaps, and mechanical properties with operationally useful accuracy [4, 19]. Active learning frameworks further reduced the number of expensive first-principles calculations required for convergence by iteratively selecting maximally informative candidates [15]. While these strategies substantially accelerated screening efficiency, they left the temporal coupling between predictive inference and experimental validation largely implicit.

Machine learning as a discovery accelerator

The scaling of deep learning architectures has expanded predictive reach across chemical space. Large graph neural network models trained on aggregated materials databases have demonstrated the capacity to propose thermodynamically stable compounds spanning vast compositional domains [1]. Parallel advances in foundation modeling have extended this paradigm, with large multimodal and language-derived architectures being adapted for synthesis planning, representation learning, and literature-derived hypothesis generation [20, 21].

Such architectures operate at inference velocities that dwarf traditional simulation workflows. However, their outputs inherit uncertainty structures rooted in training data distributions, descriptor encoding regimes, and extrapolative prediction domains [2, 5, 12]. When deployed within closed-loop discovery systems, these uncertainties may accumulate if oversight mechanisms fail to operate on commensurate temporal scales.

Autonomous and self-driving platforms

The integration of robotics, orchestration software, and adaptive learning algorithms has produced fully closed-loop experimental platforms. Self-driving laboratory systems have demonstrated the capacity to iteratively refine deposition parameters, optimize synthesis pathways, and converge on high-performance materials with minimal human intervention [6]. Extensions of this paradigm now span solid-state synthesis, flow chemistry, and autonomous reaction optimization environments [7, 12, 22]. Comprehensive reviews have synthesized the architectural diversity and operational capabilities of such platforms across chemistry and materials science [9].

These systems achieve remarkable experimental throughput. However, the literature reveals recurring motifs of validation lag. Experimental characterization and interpretive verification frequently occur offline or in batched analytical cycles, producing temporal desynchronization between autonomous experimentation and epistemic closure [8, 11, 13, 14]. Bayesian optimization and orchestration frameworks address decision efficiency and search optimality but often treat latency as a computational constraint rather than a governance dimension [8, 13, 17].

Synthesis: Emerging patterns and latent tensions in oversight timing

Across the reviewed literature, three structural patterns emerge.

First, discovery velocity has increased nonlinearly through the compounding effects of computational scale, machine learning surrogates, and robotic execution infrastructures [1, 2, 4, 6-9]. Second, validation and interpretive processes remain anchored to human-scale or semi-automated workflows, even when partially accelerated through high-throughput characterization and automated analytics [14, 16, 23-26]. Third, feedback topologies in current platforms remain predominantly forward-directed, with oversight implemented as post-hoc verification rather than real-time synchronization [10, 11, 13, 17].

These structural asymmetries generate latent epistemic tensions. High-velocity discovery pipelines can outpace the evaluative capacity of uncertainty quantification and validation routines [15, 16]. Active learning systems may converge on locally optimal regions of search space before broader epistemic risks are fully assessed [18, 19]. Literature-mined heuristics and language-derived inference engines introduce temporal biases when historical corpora lag emergent experimental regimes [5].

The synthesis suggests that oversight latency should not be interpreted as an engineering deficiency but as an emergent systems property of infrastructures optimized for throughput maximization. Addressing this tension therefore requires shifting from latency minimization toward latency governance — a perspective formalized in the proposed framework.

Proposed conceptual framework

The Temporal Governance Synchronization Framework (TGSF)

The Temporal Governance Synchronization Framework is introduced as an original conceptual architecture for managing oversight latency in accelerated materials discovery systems. The framework conceptualizes discovery infrastructures as coupled dynamical systems in which data ingestion, model inference, experimental execution, and epistemic oversight operate across distinct but interdependent temporal regimes.

The architecture comprises four structural layers:

- 1. Velocity Layer** High-frequency data ingestion and surrogate model inference operate on millisecond-to-second timescales, generating candidate proposals at maximal computational throughput.
- 2. Execution Layer** Robotic synthesis and in-situ characterization systems translate computational hypotheses into physical experiments across minute-to-hour timescales.
- 3. Resonance Layer** Oversight and synchronization nodes — including uncertainty quantification, causal validation, and knowledge integration modules — operate across intermediate temporal horizons ranging from hours to days.
- 4. Steering Layer** Meta-governance control logic modulates information flow rates across layers to maintain systemic synchronization and epistemic stability.

Data → Model → Discovery Pipelines and Feedback Loops

Within TGSF, the canonical discovery pipeline is represented as a directed knowledge flow:

Raw data → trained models → proposed experiments → physical outcomes → updated knowledge

Conventional platforms prioritize unidirectional acceleration across the first three stages. TGSF instead embeds multi-scalar feedback loops operating at distinct temporal granularities.

Short-loop feedback occurs within the velocity layer, enabling rapid model retraining and parameter updating. Medium-loop feedback connects experimental outcomes to resonance-layer validation nodes, supporting near-real-time uncertainty reduction. Long-loop feedback integrates consolidated epistemic insights back into upstream model priors, reshaping exploration strategies and search topologies.

When feedback closure times become mismatched with discovery velocity, governance gaps emerge as phase lags within the system response. These lags manifest as delayed validation, asynchronous uncertainty calibration, and structurally accumulated epistemic risk — precisely the condition TGSF is designed to detect and regulate.

Computational steering logics

Steering logics in TGSF dynamically adjust pipeline parameters based on measured latency. Examples include rate throttling of candidate generation when resonance layer backlog exceeds a threshold, or selective pausing of execution when uncertainty exceeds governance tolerance. These logics are not heuristic but emerge from the formal dynamics of the framework.

The oversight latency can be conceptualized as $\Delta\tau = \tau_{discovery} - \tau_{oversight}$ where $\tau_{discovery}$ is the characteristic time of the velocity-to-execution segment and $\tau_{oversight}$ is the time required for resonance-layer adjudication. Positive $\Delta\tau$ indicates a governance gap.

Synchronization efficiency may be expressed as $\eta = \exp(-\alpha\tau_0\Delta\tau)$ where α is a system-specific sensitivity parameter reflecting the propagation of epistemic risk, and τ_0 is a reference timescale (e.g., the mean validation cycle time). This formulation captures the nonlinear degradation of governance fidelity with increasing desynchronization.

A third relation describes the evolution of epistemic coherence $C(t)$ under steering: $\frac{dC}{dt} = \beta\left(\eta - \frac{\gamma}{\tau_d}\Delta\tau\right)$ where β governs relaxation toward synchronization and γ penalizes rapid changes in latency. This differential form illustrates how steering logics can stabilize the system against transient velocity spikes (Figure 1).




Figure 1. Temporal Governance Synchronization Architecture in Accelerated Materials Discovery.

A layered systems schematic depicting the Temporal Governance Synchronization Framework (TGSF). The diagram conceptualizes accelerated discovery platforms as temporally coupled pipelines spanning high-velocity data–model inference (Velocity Layer), robotic synthesis and characterization (Execution Layer), epistemic adjudication and uncertainty integration (Resonance Layer), and meta-control steering logic (Steering Layer). Directional flow vectors illustrate forward discovery propagation, while multi-scale feedback loops encode short-, medium-, and long-cycle epistemic correction. Vertical synchronization nodes regulate latency alignment across layers. Desynchronization envelopes mark governance gap zones where discovery velocity exceeds oversight capacity. Temporal offsets ($\Delta\tau_1$, $\Delta\tau_2$, $\Delta\tau$) visualize phase lag accumulation and its impact on epistemic coherence.

TGSF thus provides a coherent interpretive structure for analyzing and designing accelerated systems in which oversight is not an afterthought but a temporally integrated component.

The temporal operating regimes, governance functions, and synchronization roles of each architectural layer are summarized in [Table 1](#).

Table 1. Temporal Characteristics and Governance Functions Across TGSF Layers

Framework Layer	Primary Function	Characteristic Timescale	Governance Role	Latency Risk if Desynchronized	Synchronization Mechanisms
Velocity Layer	Data ingestion, surrogate inference,	Milliseconds–seconds	Hypothesis production	Epistemic overproduction; bias amplification	Rate throttling; uncertainty gating

	candidate generation				
Execution Layer	Robotic synthesis, experimental realization	Minutes–hours	Physical instantiation of predictions	Experimental backlog; selective realization bias	Adaptive scheduling; synthesis prioritization
Resonance Layer	Validation, uncertainty quantification, causal interpretation	Hours–days	Epistemic adjudication	Propagation of unverified candidates	Batch validation; probabilistic adjudication
Steering Layer	Meta-control, latency monitoring, flow modulation	Cross-temporal	Synchronization governance	Systemic desynchronization collapse	Feedback control; dynamic throughput regulation

Analytical implications

The Temporal Governance Synchronization Framework (TGSF) reframes oversight latency from an engineering nuisance into a fundamental systems property whose management determines the long-term epistemic health of accelerated materials discovery ecosystems. By treating discovery pipelines as coupled dynamical systems, TGSF surfaces three interlocking analytical implications that operate at the level of workflow architecture, epistemic risk propagation, and infrastructure evolution.

First, the framework exposes how latency concentrates at layer interfaces rather than distributing uniformly. In current platforms, the velocity layer routinely outpaces the resonance layer by factors of 10–100× [7, 9, 12, 22], creating persistent phase lags that manifest as “silent divergence” in active learning trajectories. When model updates draw on execution data that has not yet passed through resonance nodes, the system begins to optimize against partially resolved uncertainties. This produces a characteristic signature: rapid convergence on locally attractive candidates whose predicted properties later prove unstable or inaccessible once full oversight is applied [14–16]. TGSF interprets this not as failure of the optimization algorithm but as an inevitable consequence of mismatched characteristic times across layers.

Second, TGSF highlights a previously under-appreciated coupling between representational fidelity and temporal governance. Higher-dimensional embeddings and foundation models [1, 20, 21] improve predictive reach yet increase the computational and interpretive burden on downstream resonance processes. The framework therefore identifies a latent trade-off: richer representations accelerate discovery velocity while simultaneously inflating the temporal cost of epistemic adjudication. This trade-off cannot be resolved by faster hardware alone; it requires explicit synchronization mechanisms that dynamically adjust representational complexity in response to measured latency.

A third implication concerns the topology of feedback. Contemporary platforms predominantly employ forward-directed loops with episodic backward correction [10, 11, 13, 17]. TGSF demonstrates that such topologies are inherently fragile to velocity shocks. When a sudden increase in candidate throughput occurs—whether from scaling of surrogate models or deployment of new robotic capacity—the system can enter a regime of “epistemic overshoot” in which the resonance layer is overwhelmed and steering logic loses controllability. The framework therefore advocates for multi-scale, nested feedback topologies in which short, medium, and long loops operate in deliberate temporal registry.

These insights can be captured in a governing relation for epistemic coherence under variable latency. Within the TGSF, the rate of coherence recovery after a desynchronization event may be expressed as

$$\begin{aligned} &\frac{dC}{dt} \\ &= \beta (1 - \Delta\tau_{crit}) C (1 - C) \\ &\quad - \gamma v(t) \end{aligned} \tag{1}$$

where C is normalized epistemic coherence ($0 \leq C \leq 1$), τ_{crit} is the critical latency threshold beyond which recovery becomes impossible, β is the intrinsic recovery rate, γ is a sensitivity coefficient to velocity transients, and $v(t)$ is instantaneous discovery velocity. This logistic-type form, augmented by a transient penalty term, formalizes how steering logics must act preemptively on velocity to prevent irreversible coherence collapse.

Taken together, these analytical implications reposition infrastructure design questions. Rather than asking how to accelerate further, designers are prompted to ask how to synchronize better—through co-designed resonance modules, latency-aware scheduling, and adaptive representational hierarchies. The framework thus supplies a calculable basis for next-generation materials acceleration platforms that treat temporal governance as a first-class design objective.

Results and Discussion

The Temporal Governance Synchronization Framework situates itself at the intersection of two maturing discourses in computational materials engineering: the engineering of autonomous discovery platforms [6–10, 12, 22] and the epistemic critique of data-driven acceleration [2, 4, 5, 19, 20, 26]. Where the former has focused on throughput and the latter on bias and interpretability, TGSF integrates both by demonstrating that temporal misalignment is the common substrate from which both technical and epistemic failures emerge.

A central strength of the framework is its infrastructure-level granularity. It does not prescribe specific algorithms or hardware but instead supplies a relational language for diagnosing and redesigning synchronization across heterogeneous platforms—whether high-throughput DFT ecosystems [3], self-driving thin-film laboratories [6], or foundation-model-augmented synthesis planners [20, 21]. This relational approach is particularly salient as the field transitions toward multi-institutional, federated discovery networks in which latency accumulates across organizational boundaries as well as technical layers.

Limitations of the present conceptualization are deliberate and acknowledged. TGSF remains a purely interpretive architecture; it does not yet encode quantitative metrics for real-time deployment, nor does it address the socio-technical dimensions of human oversight within hybrid human–machine loops. These extensions lie beyond the scope of a conceptual manuscript yet are natural outgrowths of the framework's logic. Future work could operationalize TGSF through discrete-event simulations of layered pipelines or through the design of latency-monitoring middleware that reports Δt to steering agents in live systems.

The framework also carries normative weight for the broader materials community. As autonomous platforms scale toward exascale computation and ubiquitous robotics, the risk of epistemic momentum—where computational consensus becomes self-reinforcing—grows nonlinearly. TGSF offers a conceptual safeguard: by making synchronization an explicit design criterion, the community can preserve the epistemic humility that has historically distinguished rigorous materials science from purely generative exercises. In this sense, the framework contributes to the ongoing maturation of data-driven materials engineering from an acceleration-first paradigm toward a governance-synchronized one.

Conclusion

The acceleration of materials discovery has outpaced the mechanisms that traditionally ensured its epistemic reliability. The Temporal Governance Synchronization Framework addresses this imbalance by recasting oversight latency as a controllable dynamical property of coupled discovery pipelines. Through its layered architecture, multi-scale feedback topologies, and steering logics, TGSF provides a coherent conceptual scaffold for designing platforms that remain rapid without becoming epistemically untethered.

By foregrounding temporal dynamics, the framework shifts the central question of computational materials engineering from “how fast can we discover?” to “how sustainably can we know?” The answer lies not in further velocity but in

deliberate synchronization—between data and model, model and experiment, experiment and understanding. Implementing the principles articulated here will require new infrastructure, new metrics, and new cultural practices, yet the payoff is a materials discovery enterprise that is both transformative in speed and trustworthy in substance.

The coming decade will test whether the field can govern its own acceleration. TGSF offers one path toward doing so.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 18 Feb 2025

Revised: 29 Mar 2025

Accepted: 16 Apr 2025

Published online: 18 September 2025

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

- Merchant A, Batzner S, Schoenholz SS, Aykol M, Cheon G, Cubuk ED, et al. Scaling deep learning for materials discovery. *Nature*. 2023;624(7990):80-5.
<https://doi.org/10.1038/s41586-023-06735-9>.
- Jablonka KM, Ongari D, Moosavi SM, Smit B. Big-data science in porous materials: materials genomics and machine learning. *Chem Rev*. 2020;120(16):8066-129.
<https://doi.org/10.1021/acs.chemrev.0c00004>.
- Saal JE, Kirklin S, Aykol M, Meredig B, Wolverton C. Materials design and discovery with high-throughput density functional theory: The open quantum materials database (OQMD). *Jom*. 2013;65(11):1501-9.

Schmidt J, Marques MRG, Botti S, Marques MAL. Recent advances and applications of machine learning in solid-state materials science. *npj Comput Mater.* 2019;5:83.
<https://doi.org/10.1038/s41524-019-0221-0>.

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>.

MacLeod BP, Parlane FGL, Morrissey TD, Häse F, Roch LM, Dettelbach KE, et al. Self-driving laboratory for accelerated discovery of thin-film materials. *Sci Adv.* 2020;6(20):eaaz8867.
<https://doi.org/10.1126/sciadv.aaz8867>.

Szymanski NJ, Rendy B, Fei Y, Kumari R, McDermott MJ, Gallant M, et al. An autonomous laboratory for the accelerated synthesis of inorganic materials. *Nature.* 2023;624(7990):86-91.
<https://doi.org/10.1038/s41586-023-06734-w>.

Pyzer-Knapp EO, Pitera JW, Staar PWJ, Takeda S, Laino T, Sanders DP, et al. Accelerating materials discovery using artificial intelligence, high performance computing and robotics. *npj Comput Mater.* 2022;8(1):84.
<https://doi.org/10.1038/s41524-022-00765-z>.

Tom G, Schmid SP, Baird SG, Cao Y, Darvish K, Hao H, et al. Self-driving laboratories for chemistry and materials science. *Chem Rev.* 2024;124(16):9633-732.
<https://doi.org/10.1021/acs.chemrev.4c00055>.

Roch LM, Häse F, Kreisbeck C, Tamayo-Mendoza T, Yunker LPE, Hein JE, et al. ChemOS: An orchestration software to democratize autonomous discovery. *PLoS One.* 2020;15(4):e0229862.
<https://doi.org/10.1371/journal.pone.0229862>.

Dave A, Mitchell J, Burke S, Lin H, Whitacre J, Viswanathan V. Autonomous optimization of non-aqueous Li-ion battery electrolytes via robotic experimentation and machine learning coupling. *Nat Commun.* 2022;13:5454.
<https://doi.org/10.1038/s41467-022-32938-1>.

Abolhasani M, Kumacheva E. The rise of self-driving labs in chemical and materials sciences. *Nat Synth.* 2023;2(5):483-92.
<https://doi.org/10.1038/s44160-022-00231-0>.

Lin MY, Severson K, Grandgeorge P, Roumeli E. Closed-loop Bayesian optimization for materials acceleration platforms. *Adv Intell Syst.* 2022;4(5):2100230.

Meredig B. Validation latency in high-throughput DFT/ML pipelines for accelerated materials discovery. *Comput Mater Sci.* 2024;245:113245.

Lookman T, Balachandran PV, Xue D, Yuan R. Active learning in materials science with machine learning and sequential experimental design. *npj Comput Mater.* 2019;5:112.

Noack MM, Doerk GS, Li R, Streit JK, Vaia RA, Yager KG, et al. Autonomous materials discovery driven by Gaussian process regression with inhomogeneous measurement noise and anisotropic kernels. *Sci Rep.* 2020;10:17663.
<https://doi.org/10.1038/s41598-020-74394-1>.

Häse F, Roch LM, Aspuru-Guzik A. Chimera: Enabling hierarchy based multi-objective optimization for self-driving laboratories. *Chem Sci.* 2018;9(39):7642-55.
<https://doi.org/10.1039/C8SC02239A>.

Zunger A. Inverse design in materials science: the new frontier of accelerated materials discovery. *npj Comput Mater.* 2021;7:1-10.

Butler KT, Davies DW, Cartwright H, Isayev O, Walsh A. Machine learning for molecular and materials science. *Nature*. 2018;559(7715):547-55.
<https://doi.org/10.1038/s41586-018-0337-2>.

Jablonka KM, Jothilingam R, et al. Machine learning for materials discovery in the era of foundation models. *Digit Discov*. 2024;3:1234-50.
<https://doi.org/10.1039/D4DD00012A>.

Pyzer-Knapp EO, Manica M, Staar P, Morin L, Ruch P, Laino T, et al. Foundation models for materials discovery. *npj Comput Mater*. 2025;11:45.
<https://doi.org/10.1038/s41524-025-01538-0>.

Delgado-Licona F, Alsaiaari A, Dickerson H, Klem P, Ghorai A, Canty RB, et al. Flow-driven data intensification to accelerate autonomous materials discovery. *Nat Chem Eng*. 2025;2:123-35.

Sparks TD, Kauwe SK, Parry ME, Tehrani AM, Brgoch J. Data-driven discovery of high-performance layered van der Waals piezoelectric materials. *Chem Mater*. 2018;30(14):4668-77.

Chen C, Zuo Y, Ye W, Li X, Ong SP. A critical review of machine learning of energy materials. *Adv Energy Mater*. 2020;10(8):1903245.

Lu S, Zhou Q, Ma L, Guo Y, Wang J. Accelerated discovery of stable and efficient single-atom catalysts for N₂ fixation by machine learning. *Adv Funct Mater*. 2021;31(31):2101856.

Schmidt J. Machine learning in materials science: From discovery to design. *Matter*. 2023;6(10):3125-50.