

ORIGINAL RESEARCH

Open access

Negative Knowledge Suppression in Autonomous Materials Engineering: Archival Governance Failures

Oliver Grant^{1*}, Daniel Brooks¹, Amelia Carter²

Abstract

Autonomous materials engineering has transformed computational and data-driven discovery through self-driving laboratories, Bayesian optimization, and machine learning-guided pipelines that integrate high-throughput experimentation with predictive modeling. These systems excel at accelerating positive-outcome trajectories in materials design, from inorganic synthesis to metal-organic frameworks and functional thin films. Yet an epistemic asymmetry persists: negative knowledge—outcomes from failed reactions, suboptimal parameter spaces, unproductive compositional regions, and non-reproducible pathways—remains systematically suppressed within the archival infrastructures that underpin these ecosystems. This suppression arises not from deliberate omission but from fragmented governance mechanisms that prioritize publication-ready results, siloed data repositories, and optimization objectives indifferent to archival completeness. The present conceptual analysis synthesizes the state of autonomous experimentation, data-driven screening, and FAIR-compliant data stewardship to expose how current pipelines inadvertently amplify positive bias and erode long-term discovery efficiency. We introduce the NeGATE (Negative Epistemic Governance and Archival Transparency Ecosystem) Framework, an original systems architecture that reframes negative knowledge as an active, resonant component of the discovery loop rather than residual noise. NeGATE organizes knowledge flows across four interdependent layers—ingestion, inference, steering, and governance—while embedding computational logics that maintain traceability of suppressed signals. By foregrounding representation–inference interactions and feedback dynamics, the framework reveals infrastructure-level trade-offs that govern epistemic completeness in autonomous materials engineering. Its implications extend to the design of next-generation discovery platforms, where archival governance becomes a core computational primitive rather than a post-hoc administrative concern.

Keywords Self-driving laboratories, Autonomous materials discovery, Data-driven pipelines, Negative knowledge, Archival governance, Epistemic suppression

*Correspondence:

Oliver Grant
oliver.grant@gmail.com

¹ Department of Computational Materials Engineering, Faculty of Engineering, University of Manchester, Manchester, United Kingdom

² Department of Data-Driven Materials Science, Faculty of Engineering, University of Birmingham, Birmingham, United Kingdom

Introduction

The emergence of autonomous systems in materials discovery

The past decade has witnessed a decisive shift from manual, intuition-driven materials research toward autonomous, closed-loop platforms capable of designing, executing, and interpreting experiments with minimal human intervention. Self-driving laboratories now orchestrate robotic synthesis, in-situ characterization, and real-time machine learning feedback to navigate vast compositional and process spaces [1–3]. Landmark demonstrations, such as the autonomous

synthesis of novel inorganic materials [2] and accelerated thin-film optimization [3], illustrate how these systems compress discovery timelines from years to weeks. Complementary advances in large-language-model-guided experimentation [4] and Bayesian optimization under noisy conditions [5] have further refined the capacity of autonomous platforms to operate under uncertainty.

This transformation is grounded in the convergence of robotics, cheminformatics, and statistical learning. High-throughput screening workflows, once limited to virtual enumeration, now couple directly with physical execution environments [6, 7]. In parallel, domain-specific machine learning models have matured to predict adsorption behavior in metal-organic frameworks [8–10], zeolite formation pathways [11], and electrocatalytic performance [12], often outperforming traditional descriptor-based approaches. The literature reflects a maturing ecosystem: from early chemically-intuited screening of MOFs [13] to multi-objective Bayesian strategies in continuous-flow synthesis [14–16].

The data-driven paradigm and its epistemic foundations

Underpinning these platforms is a data-driven epistemology that treats materials space as a searchable landscape amenable to statistical inference. FAIR data principles [17] have been widely endorsed as the infrastructural backbone, yet their implementation in autonomous contexts remains uneven. Community surveys of autonomous laboratories [18] reveal that while experimental throughput has increased dramatically, the archival layer—responsible for preserving context, provenance, and outcome metadata—has not kept pace. Computational workflows increasingly rely on curated positive-result repositories, creating an implicit selection pressure that filters out negative signals before they enter model training or decision loops [19, 20].

This epistemic configuration carries structural consequences. Representation of materials systems favors high-performing candidates, while inference engines learn decision boundaries shaped predominantly by success. Steering logics, whether Gaussian-process driven [5] or reinforcement-learning informed, optimize toward reward functions that undervalue the informational content of failure. The result is a self-reinforcing cycle: experiments that deviate from expected performance are deprioritized, their data are stored in transient or inaccessible formats, and the archival record becomes progressively skewed [21, 22].

Archival governance as an emerging bottleneck

Governance failures manifest at multiple scales. At the platform level, data schemas in self-driving laboratories often lack standardized fields for capturing failure modes, experimental dead-ends, or parameter regions of low yield [23, 24]. At the community level, publication and data-sharing norms continue to reward positive outcomes, leaving negative knowledge dispersed across laboratory notebooks, unindexed repositories, or simply discarded [25, 26]. Even when negative results are retained, they are rarely indexed with the semantic richness required for machine-readable retrieval or model augmentation [27].

Recent reviews of small-data machine learning in materials science [28] and symbolic regression approaches [29] underscore the value of diverse training distributions, yet practical implementations seldom incorporate systematic negative-example curation. Studies of automated microscopy and scanning-probe workflows [30] similarly highlight the technical feasibility of comprehensive logging, yet governance protocols that would enforce archival completeness remain underdeveloped. The consequence is an accumulating epistemic debt: each autonomous campaign generates terabytes of potentially valuable negative knowledge that is effectively lost to the broader research ecosystem.

Positioning the NeGATE framework

The present work addresses this infrastructural gap through a conceptual lens. Rather than proposing incremental improvements to existing data pipelines, we articulate the NeGATE Framework as a systems-level reinterpretation of how negative knowledge can be actively stewarded within autonomous materials engineering. By foregrounding archival governance as a computational primitive, NeGATE reframes suppression not as an inevitable byproduct of optimization

but as a diagnosable and mitigable feature of current discovery architectures. The framework is developed in dialogue with the literature on self-driving laboratories [1–4, 16, 18], machine learning applications [7–14, 26–28], and data stewardship [17, 30], yet introduces an original integrative logic that treats negative knowledge as a resonant resource capable of modulating the entire discovery process.

Theoretical Background & Literature Synthesis

Advances in self-driving laboratories and autonomous experimentation

Self-driving laboratories represent the operational core of contemporary autonomous materials engineering. Early platforms demonstrated closed-loop optimization of reaction conditions in batch and flow systems [6, 14, 15], while subsequent generations integrated multi-step synthesis, in-line spectroscopy, and adaptive experimental design [1, 3, 16]. The A-Lab [2] exemplifies the capacity of such systems to discover and scale entirely new inorganic compounds, achieving synthesis targets through iterative refinement of precursor selection and processing parameters. Parallel developments in thin-film deposition [3] and battery electrolyte formulation [16] illustrate domain-specific adaptations of the same autonomous paradigm.

Large language models have recently augmented these platforms by translating natural-language objectives into executable experimental sequences [4], expanding the interface between human intent and robotic execution. Across these implementations, a common architectural pattern emerges: an experiment generator proposes conditions, a robotic executor carries them out, and a surrogate model updates its beliefs based on measured outcomes. Yet the update step is almost universally conditioned on scalar performance metrics derived from positive outcomes. Negative outcomes—reactions that yield no product, phases that decompose, or properties that fall below thresholds—are typically logged only as low-reward events and are rarely propagated as structured knowledge objects. The systemic suppression of negative knowledge manifests across multiple infrastructural strata, from ingestion schemas to inference training regimes (Table 1).

Table 1. Governance Failure Modes and Epistemic Consequences of Negative Knowledge Suppression

Governance Layer	Suppression Mechanism	Infrastructural Origin	Epistemic Consequence	Discovery Impact	NeGATE Mitigation Logic
Data ingestion	Failure data omission	Non-standard schemas	Incomplete datasets	Biased screening landscapes	Dual-stream ingestion logging
Metadata systems	Poor failure annotation	Limited ontology fields	Loss of experimental context	Reduced reproducibility	Provenance-rich tagging
Model training	Negative data underweighting	Reward-optimized datasets	Distorted decision boundaries	Overfitting to success regions	Dual-manifold embeddings
Acquisition functions	Reward asymmetry	Performance-only objectives	Failure regions ignored	Inefficient exploration	Governance-weighted utility
Archival repositories	Storage triage	Infrastructure constraints	Knowledge decay	Repeated experimental dead-ends	Persistent failure indexing
Publication ecosystems	Positive result bias	Incentive structures	Epistemic skew	Knowledge graph distortion	Failure-inclusive repositories

Platform governance	Weak archival policy	Oversight gaps	Suppression accumulation	Long-term discovery drag	Suppression coefficient tuning
---------------------	----------------------	----------------	--------------------------	--------------------------	--------------------------------

Multi-layer analysis of how suppression of failed experimental knowledge emerges across autonomous materials engineering infrastructures. The table maps suppression mechanisms to their infrastructural origins, epistemic consequences, and downstream discovery impacts, while outlining corresponding mitigation logics operationalized within the NeGATE framework.

Machine learning-driven screening and optimization in materials science

Machine learning has become the inference engine of data-driven materials research. Early demonstrations of chemically-intuited MOF screening [13] evolved into high-throughput virtual libraries coupled with graph neural networks and active learning strategies [8, 10, 12]. In zeolite synthesis, ML models now predict crystallization outcomes from precursor chemistry and process variables [11, 14]. For energy applications, active learning across intermetallic spaces has accelerated electrocatalyst discovery [12], while Bayesian optimization routines have refined flow-chemistry protocols for organic transformations [20–25].

These approaches share a methodological signature: training distributions are constructed predominantly from literature-mined positive examples, supplemented by computationally generated candidates. Negative examples, when present, are often treated as outliers or used solely for regularization rather than as carriers of structural insight. The consequence is a model landscape that is well-calibrated in high-performance regions but poorly constrained elsewhere. Small-data regimes, increasingly relevant for novel material classes [28], are particularly vulnerable to this imbalance because the informational value of each negative datum is high yet systematically underutilized.

Data management and FAIR principles in computational ecosystems

The FAIR principles—findable, accessible, interoperable, reusable—have been articulated as essential for realizing the full potential of data-driven science [17]. In materials contexts, FAIR-compliant repositories now host computed properties, experimental protocols, and characterization data. However, implementation surveys reveal persistent gaps in the treatment of negative results [18]. Metadata schemas frequently omit fields for experimental context (e.g., failed scale-up attempts, equipment-specific artifacts) that are crucial for interpreting negative outcomes. Interoperability standards focus on positive-result ontologies, leaving negative knowledge semantically orphaned.

Autonomous platforms generate heterogeneous data streams—spectra, micrographs, process logs, failure codes—yet archival systems often collapse these into summary tables optimized for publication. The result is a loss of provenance granularity that renders negative knowledge difficult to resurrect for downstream modeling or meta-analysis. Recent perspectives on data-driven materials science [27, 28] emphasize the need for richer representation of uncertainty and failure modes, yet concrete governance mechanisms that would enforce such representation remain underdeveloped.

The underrepresentation of negative results: A literature synthesis

A systematic reading of the literature reveals a consistent pattern: while technical capabilities for comprehensive data capture have advanced dramatically [1, 2, 16, 30], governance structures have not. Publications in npj Computational Materials, Nature, and Science Advances document impressive discovery rates but rarely detail the volume or nature of discarded experiments [2, 3, 12]. Reviews of machine learning in materials [26, 27] acknowledge the importance of negative examples for robust model training yet provide few examples of operationalized negative-knowledge pipelines. Studies of Bayesian optimization in chemistry [5, 15, 20] optimize acquisition functions that explicitly balance exploration and exploitation, yet the exploration term is almost always defined relative to expected positive performance rather than informational gain from negative regions.

This pattern is not accidental. It reflects deep infrastructural choices: databases are indexed by success metrics; funding and publication incentives favor positive narratives; robotic control software defaults to discarding low-yield runs to conserve storage. The epistemic cost accumulates silently. Each suppressed negative datum represents a missed opportunity to constrain model hypotheses, to map failure boundaries, and to inform future steering decisions. Over time, the collective knowledge graph of materials science becomes topologically biased—dense in high-performance clusters, sparse in the interstitial regions where transformative insights often reside.

The synthesis of these strands—autonomous execution, machine learning inference, and archival practice—points to a systemic vulnerability: negative knowledge suppression is not a peripheral data-quality issue but a core architectural feature of current discovery ecosystems. Addressing it requires more than improved metadata standards or larger repositories. It demands a reconceptualization of how negative signals are represented, propagated, and governed across the entire computational pipeline.

Proposed conceptual framework

The NeGATE framework

We propose the Negative Epistemic Governance and Archival Transparency Ecosystem (NeGATE) as an original conceptual architecture for embedding negative knowledge into the operational logic of autonomous materials engineering. NeGATE is not a software implementation or a specific algorithm but a systems-theoretic reframing that treats archival governance as an active computational layer capable of modulating data, models, and discovery trajectories in real time.

The systemic integration of suppressed experimental knowledge is formalized through the NeGATE architecture (Figure 1), which embeds archival governance directly into the computational discovery stack.




Figure 1. NeGATE architecture for negative knowledge governance.

A layered systems schematic depicting the Negative Epistemic Governance and Archival Transparency Ecosystem (NeGATE). Dual-stream ingestion channels capture both positive and negative experimental outcomes, which are processed through suppression-aware provenance filters. The augmented inference engine maintains distinct yet interacting representation manifolds, enabling failure topology encoding. Governance-weighted steering logics direct experimental exploration, while a persistent archival layer modulates epistemic resonance across the discovery pipeline. Curved feedback arcs illustrate how governed negative knowledge propagates backward to influence model training, acquisition strategies, and data retention policies.

Data → Model → Discovery pipelines

In NeGATE, the canonical discovery pipeline is restructured as a tripartite flow with embedded negative-knowledge branches. Raw experimental data enter the ingestion layer and are bifurcated. Positive data follow the conventional optimization path. Negative data are routed through a parallel representation pathway that extracts structural motifs (e.g., decomposition temperatures, incompatible precursor pairs, unstable processing windows). These motifs are not discarded but encoded as constraint manifolds within the model layer. The resulting hybrid model therefore possesses both performance-predicting and failure-predicting capabilities, enabling more nuanced steering.

Feedback loops and computational steering logics

Feedback in NeGATE is multi-scale. Short-term loops within a single campaign adjust immediate experimental parameters based on live negative signals. Medium-term loops operate across campaigns, updating global priors with archived negative knowledge. Long-term loops propagate governance insights—such as changes in suppression thresholds or indexing ontologies—back to platform design.

The steering logic itself is reformulated to treat negative knowledge as an informational asset. Rather than minimizing a scalar loss, the acquisition function seeks to maximize a composite utility that includes both expected performance gain and expected reduction in epistemic uncertainty.

This dynamic can be conceptualized as:

$$U(x) = E[f(x)] + \lambda \cdot I(x; N) \quad (1)$$

where $U(x)$ is the utility of proposing experiment x , $E[f(x)]$ is the expected positive performance, $I(x; N)$ is the mutual information between x and the negative knowledge corpus N , and λ a governance-tuned weighting factor that reflects the current state of archival completeness. The formula captures the interaction between exploitation of known high-performance regions and active exploration of failure boundaries, with λ serving as a computational dial set by the governance layer.

A second formulation formalizes the suppression dynamics across the ecosystem:

$$\sigma(t) = 1 - \frac{|N_{active}|}{|N_{total}|} \exp(-\tau t) \quad (2)$$

where $\sigma(t)$ is the instantaneous suppression coefficient at time t , $|N_{active}|$ is the volume of negative knowledge currently integrated into active models, $|N_{total}|$ is the total archived negative knowledge, and τ is a platform-specific time constant

modulated by governance refresh rate t_g . The exponential decay term reflects how timely archival intervention prevents irreversible suppression.

Finally, the resonance between layers can be expressed as:

$$\begin{aligned} \frac{dR}{dt} &= \alpha (Dp \otimes Dn) \\ &- \beta G \end{aligned} \quad (3)$$

where R is the resonance strength (a scalar measure of how effectively negative knowledge influences discovery, and Dn are positive and negative data tensors, $\otimes$ denotes a cross-manifold attention operation, G is the governance friction term, and α, β are system parameters. This equation illustrates that resonance grows when positive and negative representations interact constructively and diminishes when governance overhead becomes excessive.

Together, these formulations and the layered architecture of NeGATE provide a computational language for diagnosing and mitigating negative knowledge suppression. The framework shifts archival governance from a passive repository function to an active steering primitive, ensuring that autonomous materials engineering platforms evolve toward epistemically complete rather than merely efficient discovery.

Analytical implications

The NeGATE Framework carries direct consequences for the computational architecture of autonomous materials engineering. By elevating archival governance to a first-class computational primitive, it alters the cost–benefit surface of every layer in the discovery pipeline. Ingestion becomes more expensive in the short term—requiring richer metadata schemas and parallel negative-stream processing—but yields compounding returns in model robustness and steering precision. The Augmented Inference Engine must now maintain dual manifolds, increasing memory footprint yet reducing the frequency of catastrophic extrapolation failures when the system ventures into sparsely sampled regions of materials space.

At the level of representation–inference interactions, NeGATE reveals a previously under-appreciated trade-off. Conventional pipelines treat negative knowledge as a regularization term at best; NeGATE reframes it as a structural constraint that actively shapes the geometry of the learned manifold. This shift has measurable (though conceptual) implications for uncertainty quantification. Surrogate models trained under NeGATE governance exhibit epistemic uncertainty surfaces that are no longer radially symmetric around high-performance clusters but instead develop anisotropic ridges along documented failure boundaries. Such geometry enables more principled exploration strategies that avoid both over-exploitation of known optima and blind excursions into physically inaccessible domains.

Infrastructure-level trade-offs become explicit when governance parameters are exposed as tunable hyperparameters. Platform operators can now modulate the suppression coefficient $\sigma(t)$ [see Proposed Framework] in response to campaign objectives: high λ during early-stage materials screening to maximize informational breadth; lower λ during late-stage optimization to preserve efficiency. These choices are no longer implicit design decisions buried in code but visible, auditable policy levers. The resonance equation

$$\begin{aligned} \frac{dR}{dt} &= \alpha (Dp_{(4)} \\ &\otimes Dn) \\ &- \beta G \end{aligned}$$

Further formalizes the dynamic equilibrium between knowledge integration and governance overhead. Here, governance friction G is not a fixed constant but a function of metadata standardization effort, indexing latency, and community adoption rate. The equation implies that resonance—and therefore discovery acceleration—reaches a maximum only when governance investment is calibrated to the current state of the negative knowledge corpus. Under-investment leads

to resonance collapse; over-investment creates bureaucratic drag. NeGATE therefore supplies a computational language for negotiating these trade-offs in real time.

Epistemic risk structures are likewise transformed. Current autonomous systems accumulate “unknown unknowns” at an accelerating rate because negative signals are not systematically propagated. NeGATE converts these into “known unknowns” by maintaining an active negative knowledge ledger that is queried during every steering decision. The result is a discovery process whose risk profile is no longer dominated by silent accumulation of bias but by transparent, governable exposure to failure modes. This shift has downstream consequences for reproducibility, transferability across laboratories, and the long-term trustworthiness of machine-learned materials models.

Results and Discussion

The conceptual architecture of NeGATE sits at the intersection of three maturing strands in computational materials science: the operational maturity of self-driving laboratories [1–4, 16, 18], the increasing sophistication of data-driven inference engines [8–14, 26–28], and the still-emerging recognition that data stewardship is itself a computational problem [17, 30]. What distinguishes NeGATE is its insistence that archival governance is not ancillary infrastructure but an intrinsic component of the inference loop. This stance reframes many current challenges—poor model generalization in low-data regimes [28], unexpected phase instability in scaled synthesis [2], and the reproducibility crisis in automated workflows [18]—as symptoms of a deeper architectural omission rather than isolated technical shortcomings.

The framework also surfaces tensions that are rarely articulated in the literature. For instance, the drive toward ever-higher experimental throughput [1, 3] is frequently presented as unambiguously beneficial. Under NeGATE, throughput without corresponding governance depth becomes a liability: each additional experiment that is not archivally anchored contributes to the suppression coefficient $\sigma(t)$ and erodes long-term discovery efficiency. Similarly, the enthusiasm for large-language-model orchestration [4] must be tempered by the realization that natural-language interfaces can amplify human biases toward positive narratives unless negative knowledge is explicitly surfaced in the conversational layer.

NeGATE does not claim to resolve all tensions. It deliberately leaves open the question of optimal governance parameter regimes, acknowledging that different sub-domains—high-entropy alloys versus organic photovoltaics versus quantum materials—may require distinct resonance dynamics. The framework’s value lies instead in providing a shared conceptual vocabulary that allows these domain-specific calibrations to be discussed, compared, and iteratively refined. In this sense, NeGATE functions as a meta-infrastructure: a set of abstractions that can be instantiated differently across platforms while preserving interoperability of negative knowledge objects. Table 2 synthesizes the systemic tensions surfaced by the Negative Governance and Archival Topology Engine (NeGATE) across autonomous experimentation, inference modeling, and stewardship infrastructures. The table maps literature-anchored operational dynamics to their corresponding governance deficits and articulates the intervention logics through which NeGATE integrates archival completeness into discovery steering. Collectively, these mappings reposition negative knowledge not as residual data exhaust but as a structured inductive substrate within computational materials engineering pipelines.

Table 2. Architectural Tensions and Governance Intervention Points Identified by the NeGATE Framework

Thematic Domain	Literature Anchor(s)	Observed Systemic Tension	Governance Deficit Identified by NeGATE	Discovery Impact if Unaddressed	NeGATE Governance Intervention Logic
Self-Driving Laboratory Operations	[1–4, 16, 18]	Execution autonomy outpaces archival integration	Experimental outputs insufficiently embedded in structured negative	Reproducibility erosion; loss of tacit experimental failure states	Mandatory archival capture pipelines coupled to experimental actuation loops

			knowledge repositories		
Data-Driven Inference Engines	[8–14, 26–28]	Predictive sophistication exceeds training data representational completeness	Failure states and null results excluded from model training distributions	Poor generalization in low-data regimes; extrapolation instability	Failure-weighted dataset augmentation and suppression-aware training
Data Stewardship Infrastructure	[17, 30]	Governance treated as passive storage rather than computational layer	Absence of dynamic curation, provenance encoding, and retrieval orchestration	Knowledge fragmentation; archival opacity	Active stewardship engines integrated into inference feedback cycles
High-Throughput Experimentation Scaling	[1, 3]	Throughput expansion without proportional governance scaling	Non-anchored experimental outputs accumulate as unstructured data exhaust	Rising suppression coefficient $\sigma(t)$; declining discovery efficiency	Throughput-governance resonance calibration protocols
LLM-Mediated Laboratory Orchestration	[4]	Natural-language steering privileges positive discovery narratives	Conversational interfaces under-surface negative experimental outcomes	Bias amplification in hypothesis generation and decision steering	Conversational failure surfacing and negative knowledge prompting layers
Domain-Specific Discovery Regimes	Conceptual (framework-derived)	Uniform governance assumptions applied across heterogeneous materials classes	Misaligned archival resonance dynamics across sub-domains	Inefficient calibration of exploration–exploitation balance	Tunable governance parameterization per materials domain
Inductive Bias Formation in ML Systems	[31, 32]	Models optimized around success-centric datasets	Failure topology absent from symbolic and statistical inference	Incomplete steering logic; blind-spot persistence	Embedded failure topology encoding within model objective functions

The symbolic regression literature [32] and the cautionary analyses of machine learning pitfalls in materials discovery [31] both anticipate the need for richer inductive biases. NeGATE supplies one such bias by embedding failure topology directly into the steering logic. It thereby moves the field from a paradigm of “learning despite negative data” to one of “learning through negative data.” This transition is not incremental; it is structural. It requires platforms to treat archival completeness as a performance metric on par with experimental throughput and model accuracy.

Conclusion

Autonomous materials engineering stands at a threshold. The technical capabilities for closed-loop discovery have outpaced the epistemic infrastructure needed to steward the full spectrum of generated knowledge. Negative knowledge suppression, once an invisible background process, has become a rate-limiting factor in the long-term progress of data-

driven materials science. The NeGATE Framework offers a systems-level response: a conceptual architecture that integrates archival governance into the computational fabric of discovery itself.

By treating negative outcomes as resonant signals rather than discarded noise, NeGATE restores balance to the representation–inference–steering triad. It transforms governance from a compliance burden into a discovery accelerator. The analytical implications are concrete: richer manifolds, more principled exploration, transparent risk structures, and tunable trade-offs between efficiency and completeness. The broader implications are cultural: a research ecosystem in which failure is no longer suppressed but systematically leveraged as a source of structural insight.

The path forward lies not in adding yet another optimization algorithm or data repository, but in redesigning the very logic by which autonomous platforms remember, reason, and decide. NeGATE provides the conceptual scaffold for that redesign. Its adoption—whether in full or in modular form—will determine whether the next decade of computational materials engineering is characterized by accelerating, epistemically grounded discovery or by increasingly efficient navigation of an ever-narrower band of known successes.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 29 Jul 2024

Revised: 28 Oct 2024

Accepted: 29 Dec 2024

Published online: 18 March 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

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.

Szymanski NJ, Rendy B, Fei Y, Kumar RE, He T, Milsted D, 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>.

MacLeod BP, Parlane FGL, Morris RH, Berlinguette CP. Self-driving laboratory for accelerated discovery of thin-film materials. *Sci Adv.* 2020;6(20):eaaz8867.

Boiko DA, MacKnight R, Kline B, Gomes G. Autonomous chemical research with large language models. *Nature.* 2023;624(7992):570-8.
<https://doi.org/10.1038/s41586-023-06792-0>.

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(1):17663.

Hall BL, Taylor CJ, Labes R, Massey AF, Menzel R, Bourne RA, et al. Autonomous optimisation of a nanoparticle catalysed reduction reaction in continuous flow. *Chem Commun.* 2021;57(40):4926-9.
<https://doi.org/10.1039/D1CC00859E>.

Pardakhti M, Moharreri E, Wanik D, Suib SL, Srivastava R. Machine learning using combined structural and chemical descriptors for prediction of methane adsorption performance of metal organic frameworks (mofs). *ACS Comb Sci.* 2017;19(10):640-5.
<https://doi.org/10.1021/acscombsci.7b00056>.

Shi Z, Yang W, Deng X, Cai C, Yan Y, Liang H, et al. Machine-Learning-Assisted high-throughput computational screening of high performance metal-organic frameworks. *Mol Syst Des Eng.* 2020;5(4):725-42.
<https://doi.org/10.1039/D0ME00005A>.

Tran K, Ulissi ZW. Active learning across intermetallics to guide discovery of electrocatalysts for co2 reduction and h2 evolution. *Nat Catal.* 2018;1(9):696-703.

Zhang X, Xu Z, Wang Z, Liu H, Zhao Y, Jiang S. High-Throughput and machine learning approaches for the discovery of metal organic frameworks. *APL Mater.* 2023;11:060901.
<https://doi.org/10.1063/5.0147650>.

Borboudakis G, Stergiannakos T, Frysalis M, Klontzas E, Tsamardinos I, Froudakis GE. Chemically intuited, large-scale screening of mofs by machine learning techniques. *npj Comput Mater.* 2017;3:40.
<https://doi.org/10.1038/s41524-017-0045-8>.

Bucior BJ, Bobbitt NS, Islamoglu T, Goswami S, Gopalan A, Yildirim T, et al. Energy-Based descriptors to rapidly predict hydrogen storage in metal-organic frameworks. *Mol Syst Des Eng.* 2019;4(1):162-74.
<https://doi.org/10.1039/C8ME00050F>.

Altintas C, Altundal OF, Keskin S, Yildirim R. Machine learning meets with metal organic frameworks for gas storage and separation. *J Chem Inf Model.* 2021;61(5):2131-46.
<https://doi.org/10.1021/acs.jcim.1c00191>.

Moliner M, Román-Leshkov Y, Corma A. Machine learning applied to zeolite synthesis: The missing link for realizing high-throughput discovery. *Acc Chem Res.* 2019;52(10):2971-80.

Häse F, Roch LM, Kreisbeck C, Aspuru-Guzik A. Phoenix: A bayesian optimizer for chemistry. *ACS Cent Sci.* 2018;4(9):1134-45.
<https://doi.org/10.1021/acscentsci.8b00307>.

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(1):5454.
<https://doi.org/10.1038/s41467-022-32938-1>.

Scheffler M, Aeschlimann M, Albrecht M, Bereau T, Bungartz HJ, Felser C, et al. FAIR data enabling new horizons for materials research. *Nature.* 2022;604(7907):635-42.
<https://doi.org/10.1038/s41586-022-04501-x>.

Hung L, Yager JA, Monteverde D, Baiocchi D, Kwon H-K, Sun S, et al. Autonomous laboratories for accelerated materials discovery: A community survey and practical insights. *Digit Discov.* 2024;3(8):1273-9.

Greenaway RL, Jelfs KE. Integrating computational and experimental workflows for accelerated organic materials discovery. *Adv Mater.* 2021;33(11):2004831.
<https://doi.org/10.1002/adma.202004831>.

Eyke NS, Green WH, Jensen KF. Iterative experimental design based on active machine learning reduces the experimental burden associated with reaction screening. *React Chem Eng.* 2020;5(10):1963-72.
<https://doi.org/10.1039/D0RE00232A>.

Dunlap JH, Ethier JG, Putnam-Neeb AA, Iyer S, Luo S-X L, Feng H, et al. Continuous flow synthesis of pyridinium salts accelerated by multi-objective bayesian optimization with active learning. *Chem Sci.* 2023;14(30):8061-9.
<https://doi.org/10.1039/D3SC01303K>.

Jorayev P, Russo D, Tibbetts JD, Schweidtmann AM, Deutsch P, Bull SD, et al. Multi-objective bayesian optimisation of a two-step synthesis of p-cymene from crude sulphate turpentine. *Chem Eng Sci.* 2022;247:116938.
<https://doi.org/10.1016/j.ces.2021.116938>.

Kondo M, Sugizaki A, Khalid MI, Wathsala HDP, Ishikawa K, Hara S, et al. Energy-, time-, and labor-saving synthesis of α -ketiminophosphonates: Machine-Learning-Assisted simultaneous multiparameter screening for electrochemical oxidation. *Green Chem.* 2021;23(16):5825-31.
<https://doi.org/10.1039/D1GC01583D>.

Naito Y, Kondo M, Nakamura Y, Shida N, Ishikawa K, Washio T, et al. Bayesian optimization with constraint on passed charge for multiparameter screening of electrochemical reductive carboxylation in a flow microreactor. *Chem Commun.* 2022;58(24):3893-6.
<https://doi.org/10.1039/D2CC00124A>.

Kondo M, Wathsala HDP, Sako M, Hanatani Y, Ishikawa K, Hara S, et al. Exploration of flow reaction conditions using machine-learning for enantioselective organocatalyzed rauhut-currier and [3+2] annulation sequence. *Chem Commun.* 2020;56(8):1259-62.
<https://doi.org/10.1039/C9CC08526B>.

Kim E, Huang K, Jegelka S, Olivetti E. Virtual screening of inorganic materials synthesis parameters with deep learning. *npj Comput Mater.* 2017;3:53.
<https://doi.org/10.1038/s41524-017-0055-6>.

Choudhary K, DeCost B, Chen C, Jain A, Tavazza F, Cohn R, et al. Recent advances and applications of deep learning methods in materials science. *npj Comput Mater.* 2022;8(1):59.
<https://doi.org/10.1038/s41524-022-00734-6>.

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>.

Kalinin SV, Ziatdinov M, Hinkle J, Jesse S, Ghosh A, Kelley KP, et al. Automated and autonomous experiments in electron and scanning probe microscopy. *ACS Nano.* 2021;15(8):12653-78.

Xu P, Ji X, Li M, Lu W. Small data machine learning in materials science. *npj Comput Mater*. 2023;9:42.
<https://doi.org/10.1038/s41524-023-01000-z>.

Meredig B, Antono E, Church C, Hutchinson M, Young J, Paradiso S, et al. Can machine learning identify the next high-temperature superconductor? Identifying and avoiding common pitfalls. *npj Comput Mater*. 2018;4:21.

Wang Y, Wagner N, Rondinelli JM. Symbolic regression in materials science. *MRS Commun*. 2019;9(3):793-805.
<https://doi.org/10.1557/mrc.2019.85>.