

ORIGINAL RESEARCH

Open access

Human–AI Jurisdiction Boundaries in Computational and Data-Driven Materials Engineering

Wei Liu^{1*}, Zhang Min¹

Abstract

The computational and data-driven paradigm has fundamentally reshaped materials engineering, enabling the navigation of vast chemical and structural spaces through machine learning, high-throughput computation, and autonomous workflows. Yet this transformation has also exposed a critical conceptual gap: the absence of explicit, structured boundaries that govern the division of cognitive labor between human experts and artificial intelligence systems. Without such boundaries, AI contributions risk overstepping domains requiring physical intuition, ethical judgment, and contextual synthesis, while human oversight may inadvertently constrain the scale and speed that define modern discovery pipelines. This manuscript introduces the Epistemic Jurisdiction Framework (EJF), an original systems-level model that delineates jurisdiction layers, interfaces, and feedback mechanisms tailored to the materials discovery ecosystem. The EJF maps the flow from raw data to validated discovery through distinct zones of human primacy, AI autonomy, and negotiated hybrid spaces, emphasizing representation–inference interactions and computational steering logics. Grounded in the recent literature on machine learning for materials, explainable systems, and data-driven infrastructures, the framework offers a conceptual scaffold for infrastructure-level design rather than performance optimization. Its implications extend to the construction of more robust, interpretable, and sustainable computational ecosystems in which human and AI capabilities are aligned rather than blurred. The EJF thereby provides a foundation for next-generation materials engineering platforms that preserve epistemic integrity while fully exploiting computational scale.

Keywords Computational materials science, Human–AI collaboration, Jurisdiction boundaries, Epistemic frameworks, Data-driven materials discovery, AI governance

*Correspondence:

Wei Liu

wei.liu@gmail.com

¹ Department of Materials Informatics, School of Materials Science, Shanghai Jiao Tong University, Shanghai, China

Introduction

The emergence of data-driven paradigms in materials engineering

Over the past decade, materials engineering has transitioned from a predominantly experiment- and theory-driven discipline to one increasingly orchestrated by computational and data-centric methods. High-throughput density-functional theory calculations, automated synthesis platforms, and large-scale materials databases have

collectively expanded the searchable space of possible compounds by orders of magnitude [1–4]. Machine learning models now routinely predict formation energies, band gaps, and mechanical properties directly from compositional or structural descriptors, often surpassing traditional physics-based approaches in speed and scope [5–8]. Generative models further enable inverse design, proposing candidate structures that satisfy target property profiles [3, 9].

This shift has been enabled by the convergence of three infrastructures: (i) massive materials repositories that serve as training corpora [4, 10, 11], (ii) graph-based and transformer architectures that capture chemical and physical invariances [8, 12], and (iii) active-learning and Bayesian optimization routines that close the loop between prediction and validation [13, 14]. The result is a discovery pipeline that operates at a scale and velocity unattainable by human effort alone. Yet this very capability has surfaced a structural tension. AI systems excel at pattern extraction across high-dimensional spaces but lack the embodied, contextual, and normative reasoning that human experts bring to bear when interpreting anomalies, assessing synthesizability, or navigating regulatory and ethical constraints [15-18].

The blurring of human and machine roles

Contemporary literature reveals a progressive erosion of clear role distinctions. Early workflows positioned machine learning as a screening tool whose outputs were filtered by human experts [19, 20]. More recent systems embed AI deeper into the pipeline, delegating not only prediction but also experiment selection, hypothesis refinement, and even literature interpretation [3, 10, 13]. Autonomous laboratories now operate with minimal human intervention, raising questions about accountability when a predicted material fails synthesis or exhibits unexpected toxicity [11, 14].

This blurring is not merely operational; it is epistemic. When an AI model identifies a novel perovskite composition with promising photovoltaic properties, the attribution of “discovery” becomes ambiguous. Is credit due to the human who curated the training set, the architect of the graph network, the operator of the robotic platform, or the model itself [16, 17, 21]? The literature on explainable machine learning in materials has responded by developing post-hoc interpretability tools and uncertainty quantification methods [15-17], yet these approaches treat interpretability as a retrofit rather than a structural design principle.

The need for explicit jurisdiction boundaries

The absence of formalized boundaries creates several systemic vulnerabilities. First, epistemic risk accumulates when AI outputs are accepted in domains where human judgment is irreplaceable—such as the assessment of

long-term stability under real-world conditions or the integration of domain knowledge that is sparsely represented in training data [15, 18]. Second, over-reliance on AI can suppress serendipitous human insight that arises from analogical reasoning across disparate materials classes [2, 4]. Third, the lack of clear interfaces complicates the design of collaborative platforms, leading to inefficient hand-offs and duplicated effort [11, 14].

These challenges are infrastructure-level rather than algorithmic. They concern how computational ecosystems are architected to distribute authority, responsibility, and cognitive load. The materials science community has begun to recognize this through discussions of “human-in-the-loop” systems and “co-pilot” metaphors [16-18], yet the field still lacks a unifying conceptual model that treats jurisdiction as a first-class design variable.

Positioning the current contribution

This manuscript addresses the gap by proposing the Epistemic Jurisdiction Framework (EJF)—a novel conceptual architecture that explicitly delineates, negotiates, and maintains human–AI boundaries within computational and data-driven materials engineering workflows. Unlike existing taxonomies that classify tasks as human or machine, the EJF models jurisdictions as dynamic, layered, and interdependent structures with defined interfaces and feedback mechanisms. The framework is developed entirely at the conceptual and systems level, drawing integrative insights from the literature to illuminate how boundaries can be engineered to enhance rather than constrain discovery capacity.

Theoretical Background & Literature Synthesis

Machine learning foundations in materials science

The foundational literature from 2017 onward established machine learning as a core instrument in materials engineering. Schmidt and colleagues [1] surveyed the rapid adoption of supervised and unsupervised methods for property prediction and materials classification, highlighting the transition from descriptor-based models to end-to-end learning on crystal graphs. Butler *et al.* [2] provided a broader perspective on the potential of machine learning to accelerate molecular and materials discovery, emphasizing

the complementarity of data-driven approaches with first-principles computation. These works collectively demonstrated that the materials community had moved beyond proof-of-concept studies to the construction of production-scale prediction engines [5-7, 9].

Subsequent contributions refined the methodological toolkit. Graph neural networks emerged as a dominant paradigm for capturing many-body interactions in crystalline and amorphous systems [8, 12]. Similar architectures showed that even elemental composition alone could yield surprisingly accurate property predictions, underscoring the power of data-driven representations. At the same time, large-scale benchmarking efforts [6, 7] revealed persistent gaps between model performance on in-distribution data and generalization to novel chemical spaces, pointing to the enduring necessity of human-guided feature engineering and validation.

Explainability and interpretability as boundary signals

A parallel thread in the literature has focused on making AI outputs legible to human experts [15]. Some studies reviewed explainable machine learning techniques tailored to materials problems, ranging from attention mechanisms in graph networks to symbolic regression approaches [16, 21]. Another study argued that explainability is not merely a usability feature but a prerequisite for trustworthy deployment in engineering contexts [17]. These contributions implicitly acknowledge that interpretability functions as a boundary mechanism: it allows human experts to reclaim jurisdiction when model explanations reveal inconsistencies with physical principles or domain knowledge.

Yet the literature also reveals limitations. Most explainability methods operate post hoc and are applied to models that were trained without explicit consideration of human oversight requirements [15]. This creates a structural misalignment: the model optimizes for predictive accuracy while human experts must later reconstruct the reasoning path. The EIJF addresses this by embedding interpretability interfaces directly into the jurisdictional architecture.

High-throughput and autonomous discovery systems

The scaling of computational discovery has been documented in several landmark studies [3]. demonstrated the use of deep learning to screen millions of candidate materials, identifying novel compounds that were subsequently synthesized. A study showed how unsupervised word embeddings extracted from the scientific literature could capture latent materials knowledge, enabling zero-shot property prediction [10]. Active learning strategies [13] further closed the loop by using model uncertainty to select the most informative experiments, effectively delegating exploration strategy to the machine.

These advances have pushed the field toward closed-loop, autonomous platforms [11, 14]. However, the same literature repeatedly notes that full autonomy remains elusive because human judgment is required at critical junctures: defining meaningful objectives, interpreting unexpected results, and deciding when to terminate a search. The synthesis therefore reveals a recurring pattern: AI systems expand the frontier of what is computationally feasible, yet human expertise defines the boundaries of what is scientifically and practically meaningful.

Data ecosystems and infrastructure perspectives

A growing body of work has examined the infrastructural underpinnings of data-driven materials science [4, 11] described the emergence of materials intelligence ecosystems that integrate databases, ontologies, and analysis tools. A study outlined both the opportunities and the governance challenges inherent in large-scale data sharing [14]. These contributions shift the focus from individual algorithms to the socio-technical systems that sustain discovery at scale.

Within these ecosystems, the division of labor is rarely made explicit. Data curation, ontology development, and quality assurance remain predominantly human activities [4, 18], while model training and inference are increasingly automated. The literature thus provides rich empirical grounding for the need to formalize these divisions as jurisdictional layers rather than ad hoc arrangements.

Human expertise in the data-driven loop

Several studies have begun to articulate the irreplaceable role of human experts [16, 18] emphasized that human intuition remains essential for hypothesis generation, anomaly detection, and the integration of multimodal knowledge. Recent studies noted that even the most sophisticated models benefit from human-defined constraints that encode physical laws or application-specific requirements [20, 22, 23].

This body of work collectively supports the central thesis of the present framework: optimal discovery emerges not from maximizing AI autonomy but from the deliberate engineering of boundaries that allow each agent—human and machine—to operate in its zone of comparative advantage while maintaining continuous, structured interaction.

Proposed conceptual framework

The Epistemic Jurisdiction Framework (EJF)

The Epistemic Jurisdiction Framework (EJF) conceptualizes the materials discovery pipeline as a layered architecture in which jurisdiction—the right and responsibility to make specific classes of decisions—is explicitly allocated, negotiated, and maintained. Unlike task-allocation models that treat human and AI roles as static, the EJF treats jurisdictions as dynamic surfaces defined by epistemic risk, computational scale, and interface protocols.

The framework comprises four primary layers, each characterized by a dominant jurisdiction type:

- 1. Data Stewardship Layer (Human Primacy):** Human experts define data provenance, quality thresholds, ontology structures, and ethical constraints. AI assists through anomaly detection and metadata enrichment but cannot override human-defined governance rules.
- 2. Representation Learning Layer (Hybrid with AI Lead):** AI constructs latent representations from curated data, while humans set representational priors (e.g., symmetry constraints, physically motivated descriptors) and perform periodic audits of representation fidelity.
- 3. Inference and Prediction Layer (AI Primacy with Human Veto):** AI performs high-throughput inference and uncertainty quantification. Human jurisdiction is exercised through predefined veto thresholds based on epistemic risk profiles.

4. Interpretation and Discovery Steering Layer

(Human Primacy with AI Augmentation): Humans synthesize predictions into scientific narratives, assess synthesizability, and steer future exploration. AI proposes alternative interpretations and exploration trajectories for human consideration.

These layers are connected by standardized interfaces that encode the transfer of authority. Each interface includes (i) a data contract specifying required provenance and uncertainty metadata, (ii) an epistemic checkpoint where human review is mandatory, and (iii) a steering protocol that allows higher layers to modulate lower-layer behavior.

Pipelines, feedback loops, and steering logics

The primary discovery pipeline flows unidirectionally from data stewardship through representation, inference, and finally to steering. However, the EJF incorporates three classes of feedback loops:

- **Performance feedback:** Quantitative metrics from downstream validation are propagated upward to refine representations and data curation rules.
- **Epistemic feedback:** Human experts inject qualitative judgments (e.g., “this prediction contradicts established phase stability trends”) that trigger re-training or constraint addition.
- **Exploration feedback:** AI-proposed search directions are evaluated by humans and either accepted, modified, or rejected, thereby shaping subsequent iterations.

Computational steering logics operate at each interface. In the representation layer, humans can inject soft constraints via regularization terms or by augmenting the loss landscape. In the inference layer, steering manifests as adaptive sampling weights that favor regions of high human interest. In the steering layer, AI acts as a recommendation engine whose proposals are filtered through human-defined utility functions. Operating as a vertically integrated jurisdiction stack with dynamic feedback coupling across epistemic layers (**Figure 1**).

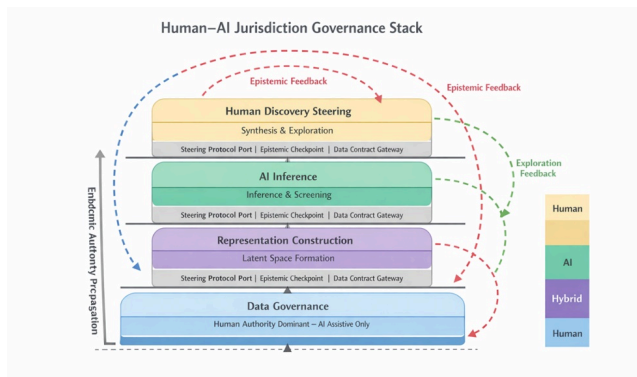


Figure 1. Epistemic Jurisdiction Framework (EJF) for computational and data-driven materials engineering.

The framework conceptualizes materials discovery as a vertically stratified jurisdiction architecture spanning four governance layers: Data Stewardship, Representation Learning, Inference & Prediction, and Interpretation & Discovery Steering. Each layer encodes a dominant locus of epistemic authority while preserving hybrid negotiation zones and veto-enabled interfaces. Horizontal boundary bands formalize authority transfer through data contracts, epistemic checkpoints, and steering protocols. Solid vertical arrows represent the primary discovery pipeline, whereas dashed curved arcs denote performance, epistemic, and exploration feedback loops. The architecture reframes human–AI collaboration as a governed systems structure in which jurisdiction boundaries function as active control surfaces for aligning computational scale with epistemic integrity.

Conceptual formalizations

Within the EJF, several dynamics are captured through abstract symbolic relations.

The jurisdiction transfer at the data-to-representation interface may be expressed as: $T_{\{dr\}} = \frac{A_r}{H_s + A_r}$ where A_r denotes the representational capacity contributed by AI and H_s the stewardship depth provided by human experts. This ratio formalizes the gradual shift from human-controlled data governance to AI-driven pattern extraction while ensuring that T_{dr} remains bounded by human stewardship.

The strength of epistemic feedback across layers can be conceptualized as: $F_e = \delta \cdot (U_h - U_a)$ where U_h is the

uncertainty perceived by the human expert, U_a is the uncertainty reported by the AI system, and δ is an interface-specific sensitivity parameter. Positive values of F_e trigger upward propagation of corrective signals, implementing a proportional steering mechanism.

Finally, the overall coherence of the discovery process is captured by the boundary alignment metric:

$$B = \prod_{l=1}^4 (1 - |J_l - 0.5|) \quad (1)$$

where J_l is the normalized jurisdiction index (0 = full human control, 1 = full AI control) in layer l . The product favors configurations in which no single layer is dominated by one agent, thereby promoting balanced human–AI symbiosis.

These formalizations are not predictive equations but interpretive devices that make visible the trade-offs and interactions engineered into the EJF can be seen in [Table 1](#).

Table 1. Jurisdiction Allocation Matrix Across Human–AI Collaborative Layers in the Epistemic Jurisdiction Framework

Framework Layer	Dominant Jurisdiction Agent	AI Functional Authority	Human Functional Authority
Data Stewardship	Human Primacy	Metadata parsing, anomaly detection	Provenance, ontology, structural, ethical, constraints
Representation Learning	Hybrid (AI Lead)	Latent encoding, feature abstraction	Description, symmetry, constraints, representation, accuracy
Inference & Prediction	AI Primacy	Property prediction, generative screening, uncertainty modeling	Thresholds, interpretation, review

Interpretation & Steering	Human Primacy	Trajectory recommendation, alternative hypothesis generation	Synthes reasc dom trans discc direc
---------------------------	---------------	--	-------------------------------------

Analytical implications

The Epistemic Jurisdiction Framework (EJF) reframes the materials discovery pipeline as a governed information-processing system in which jurisdiction boundaries function as active control surfaces rather than passive hand-off points. This perspective yields several systems-level insights that reshape how computational infrastructures are conceptualized and engineered.

At the most fundamental level, the EJF reveals that epistemic risk is not distributed uniformly across the pipeline but concentrates at jurisdiction interfaces. When human stewardship in the data layer is insufficiently coupled to AI-led representation learning, representational drift can propagate undetected, producing inference outputs that appear statistically robust yet violate latent physical constraints. The framework therefore implies that infrastructure designers must treat interface protocols as primary risk-mitigation instruments, embedding mandatory provenance metadata and human-auditable constraint layers at every transition.

A second implication concerns the dynamics of discovery velocity. Traditional data-driven workflows optimize for throughput by minimizing human intervention; the EJF demonstrates that strategic retention of human jurisdiction at the interpretation and steering layer can paradoxically accelerate long-term progress. By constraining AI exploration to regions of high epistemic coherence, the framework reduces the fraction of computational effort expended on scientifically sterile regions of chemical space. This can be formalized through the boundary-coherence relation introduced earlier, extended here to include temporal evolution:

The cumulative discovery efficiency over successive pipeline cycles may be conceptualized as

$$\eta(t) = B(t) \cdot \exp\left(-\lambda \int_0^t |Fe(\tau)| d\tau\right)$$

where $B(t)$ is the instantaneous boundary alignment, λ is a damping coefficient representing the cost of unresolved epistemic feedback, and the integral accumulates corrective

interventions. The expression captures the insight that well-maintained jurisdictions produce compounding returns: each resolved feedback loop tightens the coherence envelope, thereby increasing the fraction of generated candidates that survive downstream validation.

A third analytical consequence emerges in the representation–inference interaction. The EJF positions representation learning not as a preprocessing step but as a jurisdictionally contested domain where human priors and AI pattern extraction negotiate a shared latent space. When this negotiation is explicit, the resulting representations become more robust to distributional shift—an outcome observed across multiple materials classes in the literature [5, 8, 12]. The framework therefore suggests that next-generation platforms should expose representation-audit interfaces that allow human experts to inject symmetry-breaking terms or physically grounded regularizers directly into the training objective, transforming what has been an opaque optimization process into a transparent, co-authored representational contract.

Finally, the EJF illuminates infrastructure trade-offs at the ecosystem scale. Allocating excessive jurisdiction to AI in the inference layer maximizes short-term screening volume but erodes the human capacity for analogical transfer across materials domains—an capacity repeatedly highlighted as essential for paradigm-shifting discoveries [2, 4, 14]. Conversely, overly conservative human veto thresholds can throttle the generative potential of modern foundation models [3]. The framework advocates for adaptive jurisdiction surfaces whose allocation parameters are themselves subject to higher-order steering, creating a meta-control layer that continuously recalibrates the human–AI division based on the maturity of the materials subdomain under investigation.

These implications collectively shift the design objective from maximizing model accuracy to optimizing jurisdictional coherence, thereby aligning computational scale with epistemic integrity.

Results and Discussion

The Epistemic Jurisdiction Framework addresses a structural lacuna that has persisted across the data-driven materials literature: the treatment of human and machine contributions as interchangeable resources rather than as agents possessing fundamentally distinct epistemic

affordances. By making jurisdiction an explicit architectural primitive, the EJF moves beyond the “human-in-the-loop” rhetoric that has dominated recent reviews [16–18] and offers instead a systems-theoretic foundation for platform design.

This contribution is particularly timely given the rapid integration of large language models and autonomous laboratory platforms into materials workflows [3, 11]. Such systems amplify both the opportunities and the hazards of blurred boundaries. The framework’s layered architecture provides a scaffold for embedding governance mechanisms at the point of deployment rather than retrofitting them after deployment—an approach that aligns with emerging calls for responsible AI in scientific infrastructure [14].

Several design principles follow directly from the EJF. First, platform architectures should expose jurisdiction dashboards that visualize current boundary states, epistemic feedback flux, and projected risk accumulation. Second, data ontologies must be co-evolved with jurisdictional protocols so that provenance metadata carries explicit authority signatures. Third, training objectives for representation models should incorporate boundary-alignment penalties that reward configurations in which human and AI contributions remain distinguishable yet synergistic.

The framework also carries implications for education and community practice. Training the next generation of computational materials scientists must include explicit instruction in jurisdiction engineering—teaching researchers to reason not only about model hyperparameters but also about the allocation of decision rights within socio-technical pipelines. This curricular shift would mirror the infrastructure-level perspective already evident in leading materials data initiatives [4, 11].

Limitations of the present work are conceptual rather than empirical. The EJF is deliberately abstract, providing a high-level map rather than implementation blueprints. Translating its principles into production platforms will require domain-specific refinements, particularly in subfields such as polymer informatics or high-entropy alloys where data sparsity and physical complexity differ markedly. Nevertheless, the framework’s modularity ensures that its core jurisdictional logic remains portable across these contexts.

Future extensions could incorporate multi-agent jurisdictions in which multiple AI systems with complementary architectures negotiate among themselves under human oversight, or explore the integration of ethical and societal constraints as additional jurisdictional layers. The central insight, however, remains invariant: sustainable progress in computational materials engineering will depend less on ever-larger models and more on the deliberate engineering of the cognitive boundaries that connect human insight with machine scale.

Conclusion

The computational and data-driven transformation of materials engineering has delivered unprecedented capacity to explore chemical space, yet it has simultaneously created an urgent need for structured governance of human–AI interaction. The Epistemic Jurisdiction Framework introduced here offers a conceptual architecture that makes these interactions explicit, negotiable, and maintainable. By delineating layered jurisdictions, standardized interfaces, and feedback-driven steering logics, the EJF provides a foundation for discovery ecosystems that preserve epistemic integrity while fully exploiting computational scale.

Rather than viewing human expertise as a bottleneck to be minimized, the framework positions it as a critical regulatory substrate that channels AI capabilities toward scientifically consequential outcomes. This reorientation has direct consequences for infrastructure design, platform governance, and scientific training. As materials discovery platforms continue to evolve toward greater autonomy, the principles articulated in the EJF can serve as design invariants that safeguard the human role in knowledge creation.

The ultimate promise of the framework is a materials engineering discipline in which human and artificial intelligence operate as complementary epistemic agents within a coherently bounded system. In such an ecosystem, the boundaries between human and machine are not barriers to progress but the very structures that enable it.

Acknowledgements

None

Conflict of interest

None

Ethics statement

None

Financial support

None

Received: 23 Oct 2024 Revised: 02 Dec 2024 Accepted: 16 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

- 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.
- Butler KT, Davies DW, Cartwright H, Isayev O, Walsh A. Machine learning for molecular and materials science. *Nature.* 2018;559(7715):547-55.
- Merchant A, Batzner S, Schoenholz SS, Aykol M, Cheon G, Cubuk ED. Scaling deep learning for materials discovery. *Nature.* 2023;624(7990):80-5.
- 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>.
- Xu P, Ji X, Li M, Lu W. Small data machine learning in materials science. *npj Comput Mater.* 2023;9(1):42. <https://doi.org/10.1038/s41524-023-01000-z>.
- 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>.
- 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. <https://doi.org/10.1038/s41524-020-00406-3>.
- 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>.
- Choubisa H, Todorović P, Pina JM, Parmar DH, Li Z, Voznyy O, et al. Interpretable discovery of semiconductors with machine learning. *npj Comput Mater.* 2023;9(1):117. <https://doi.org/10.1038/s41524-023-01066-9>.
- 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>.
- Batra R, Song L, Ramprasad R. Emerging materials intelligence ecosystems. *Matter.* 2020;3(5):1433-8.
- 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.
- 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.
- Morgan D, Jacobs R. Opportunities and challenges for machine learning in materials science. *Annu Rev Mater Res.* 2020;50:71-103.
- Zhong X, Gallagher B, Liu S, Kailkhura B, Hiszpanski A, Han TY-J. Explainable machine learning in materials science. *npj*

Comput Mater. 2022;8:204.

<https://doi.org/10.1038/s41524-022-00884-7>.

Oviedo F, Lavista Ferres J, Buonassisi T, Butler KT.

Interpretable and Explainable Machine Learning for Materials Science and Chemistry. *Acc Mater Res.* 2022;3(6):597-607.

Pilania G. Machine learning in materials science: From explainable predictions to autonomous design. *Comput Mater Sci.* 2021;193:110360.

<https://doi.org/10.1016/j.commatsci.2021.110360>.

Wang A, Murdock R, Kauwe S, Oliynyk A, Gurlo A, Brgoch J, et al. Machine learning for materials scientists: an introductory guide. *J Mater Res.* 2020;35(24):4077-85.

Liu Y, Zhao T, Ju W, Shi S. Materials discovery and design using machine learning. *J Materomics.* 2017;3(3):159-77.

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.

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>.

Choudhary K, DeCost B, Tavazza F. Machine learning for materials science: Recent progress and emerging applications. *J Phys Condens Matter.* 2022;34(4):043001.

Choudhary K, DeCost B, Chen C, Jain A, Tavazza F, Cohn R, et al.. Recent advances and applications of machine learning in solid-state materials science. *npj Comput Mater.* 2022;8(1):59.

<https://doi.org/10.1038/s41524-022-00734-6>.