ORIGINAL RESEARCH Open access

The Ethics of Materials Acceleration: A Conceptual Analysis of Sustainability, Labor, and Dual-Use Tensions

Claire Dupont[1*], Julien Martin[1], Sophie Bernard[2]

Abstract

Materials acceleration—the compression of materials discovery timelines through automated experimentation and data-driven decision loops in self-driving laboratories (SDLs) and materials acceleration platforms (MAPs)—is reshaping contemporary materials science. While widely promoted for its efficiency and sustainability potential, accelerated discovery also introduces ethical tensions that remain insufficiently theorized. This conceptual paper develops a novel framework to analyze how acceleration restructures ethical challenges across five interdependent dimensions: sustainability, labor, dual-use risk, inequity, and governance. Drawing exclusively on peer-reviewed literature published, the analysis shows that compressed timelines and autonomous decision loops function as ethical multipliers, intensifying trade-offs rather than resolving them. Claimed computational and infrastructural burdens often offset sustainability gains; automation reconfigures scientific labor and risks epistemic deskilling; accelerated optimization amplifies dual-use vulnerabilities; access asymmetries widen global research inequities; and existing governance mechanisms lag behind acceleration velocity. To integrate these dynamics, the paper introduces the ethical acceleration tension matrix. This multidimensional framework models feedback interactions and identifies leverage points for ethical steering under conditions of speed and autonomy. By foregrounding interdependence, feedback velocity, and equilibrium steering—without recourse to empirical data—this work provides a foundational conceptual logic for responsible acceleration in applied artificial intelligence for materials science. Implications are outlined for platform design, governance, and education to align innovation velocity with societal safeguards.

Keywords Materials acceleration, Self-driving laboratories, AI ethics, Sustainability trade-offs, Labor deskilling, Dual-use risk

*Correspondence:
Claire Dupont
claire.dupont@outlook.com
[1] Department of Materials Data Analytics, Faculty of Engineering, University of Bordeaux, Bordeaux, France
[2] Department of Artificial Intelligence Systems, Faculty of Engineering, University of Toulouse, Toulouse, France

Introduction

The integration of artificial intelligence (AI), robotics, and automated experimentation into materials science has inaugurated a qualitatively new mode of scientific discovery characterized by compressed timelines, closed-loop decision-making, and reduced reliance on continuous human intervention. Within this paradigm, materials acceleration platforms (MAPs) and self-driving laboratories (SDLs) combine robotic synthesis, high-throughput characterization, machine learning (ML)–based hypothesis generation, and adaptive optimization into autonomous experimental pipelines capable of iteratively designing, executing, analyzing, and refining experiments with minimal human input [1, 2]. These systems do not merely accelerate traditional workflows; rather, they reconfigure the epistemic structure of materials research by shifting from sequential, hypothesis-driven experimentation toward parallelized, probabilistic, and continuously adaptive exploration [3, 4].

This phenomenon—hereafter referred to as *materials acceleration*—has been widely framed as a transformative advance for addressing pressing societal challenges. Proponents argue that MAPs and SDLs can dramatically shorten the path from concept to validated material, enabling rapid development of next-generation materials for clean energy technologies, carbon capture, sustainable electronics, and advanced manufacturing [5]. In this narrative, acceleration is positioned as both a technical and moral good: faster discovery promises earlier deployment of sustainable solutions in the face of climate urgency and resource constraints.

Yet the ethical implications of compressing discovery timelines through autonomous systems remain underexamined. While technical performance metrics—such as optimization speed, prediction accuracy, or experimental throughput—dominate current evaluations of SDLs, acceleration introduces a distinct class of ethical tensions that cannot be reduced to conventional AI risk discussions. These tensions arise precisely because materials acceleration operates at the intersection of physical experimentation, resource-intensive computation, high-stakes applications, and shifting scientific labor practices. As acceleration increases in speed and scale, ethical considerations become tightly coupled rather than separable, demanding conceptual scrutiny beyond case-by-case assessment.

This paper focuses on five interrelated ethical axes that recur across the emerging literature: sustainability claims, labor displacement and deskilling, dual-use risks, global inequity, and governance deficits. Sustainability is frequently invoked as a core justification for materials acceleration, with claims that autonomous optimization reduces waste, minimizes failed experiments, and enables targeted discovery of environmentally benign materials [5]. However, such claims often abstract away the substantial computational, infrastructural, and hardware footprints of AI-driven platforms, including energy- and water-intensive model training, reliance on critical minerals for robotic and computing hardware, and the rebound effects of faster innovation cycles [6]. As a result, acceleration risks producing a paradox in which purported sustainability gains are offset—or even negated—by upstream and downstream resource demands. The five ethical axes examined in this study and their associated tensions in accelerated materials discovery are summarized in **Table 1**.

Table 1. Core ethical axes and tensions in materials acceleration

Ethical axis	Conceptual definition	Primary ethical tension under acceleration	Representative manifestation in SDLs/MAPs
Sustainability (S)	Claims that accelerated discovery reduces environmental impact through efficiency and targeted optimization	Procedural efficiency versus system-level resource consumption	Reduced experimental waste offset by energy-intensive ML training, robotic infrastructure, and critical material dependencies
Labor (L)	Transformation of scientific work through automation and autonomy	Productivity gains versus deskilling and epistemic concentration	Displacement of technicians and early-career researchers; erosion of tacit experimental knowledge
Dual-Use (D)	Potential for materials to serve both beneficial and harmful applications	Sustainability-driven optimization versus misuse potential	Battery, alloy, or semiconductor advances enabling both clean energy and military systems
Inequity (I)	Unequal access to accelerated discovery infrastructure and influence	Innovation concentration versus global research inclusion	Capital-intensive SDLs are concentrated in Global North institutions
Governance (G)	Ethical oversight mechanisms shaping objectives, constraints, and accountability	Innovation velocity versus regulatory and ethical lag	Absence of domain-specific protocols for autonomous materials experimentation

Labor dynamics constitute a second axis of ethical tension. Automation within SDLs redistributes scientific labor by relegating routine synthesis and characterization to machines while elevating human roles toward system design, supervision, and strategic decision-making [7]. Although this shift is frequently framed as augmentation rather than replacement, concerns arise about deskilling, the displacement of technicians and early-career researchers, and the erosion of tacit, hands-on knowledge traditionally acquired through iterative experimentation. Over time, such epistemic deskilling may reshape the culture of materials science itself, concentrating authority among those who control algorithms and infrastructure while narrowing experiential learning pathways.

A third axis involves dual-use risks inherent to accelerated materials discovery. Materials optimized for sustainability-oriented applications—such as high-performance batteries, lightweight alloys, or advanced semiconductors—may simultaneously enable military, surveillance, or other harmful uses [8]. Acceleration amplifies these risks by rapidly proliferating materials knowledge through open datasets, automated pipelines, and scalable optimization, potentially outpacing ethical review mechanisms and export-control regimes. The blurring of boundaries between benign and harmful applications complicates responsibility attribution and challenges existing oversight frameworks.

Inequity forms a fourth axis, arising from the capital-intensive nature of MAPs and SDLs. The deployment of autonomous laboratories requires substantial investments in robotics, computation, data infrastructure, and skilled personnel, advantages disproportionately concentrated in elite institutions and well-resourced regions of the Global North [9, 10]. As a result, materials acceleration risks widening global disparities in discovery capacity, reinforcing dependencies on imported advanced materials, and marginalizing under-resourced researchers from shaping research agendas. These inequities are not merely distributive but epistemic, influencing which problems are prioritized and which values are encoded into autonomous systems.

Finally, governance emerges as a cross-cutting challenge. Despite rapid advances in autonomous experimentation, there remains a lack of domain-specific ethical protocols tailored to accelerated materials research. Existing AI ethics frameworks—emphasizing principles such as transparency, fairness, or accountability—provide important foundations but were largely developed for software-centric systems and decision-making contexts. They inadequately address the distinctive features of materials acceleration: physical risk, resource materiality, dual-use propagation, and the temporal compression that shortens windows for ethical intervention. Governance deficits are further exacerbated by uneven regulatory landscapes, fragmented oversight, and misalignment between innovation velocity and policy response times.

This manuscript offers a deliberately **conceptual** analysis of these challenges. It introduces no new empirical data, simulations, or experimental results. Instead, it develops an original theoretical framework through a systematic synthesis of peer-reviewed literature. The central contribution is a novel conceptual model that treats acceleration not as a neutral efficiency gain, but as an *ethical multiplier*. This force amplifies tensions across sustainability, labor, dual-use, inequity, and governance through feedback dynamics and shortened response horizons. The proposed framework conceptualizes these interdependencies as a structured tension matrix, moving beyond static ethical checklists toward a logic of interaction, velocity, and equilibrium steering [11, 12].

Historically, materials acceleration builds upon decades of automation and high-throughput experimentation, but acquires new potency through ML-orchestrated closed loops. Early SDL implementations demonstrated rapid optimization of thin-film properties using robotic synthesis coupled with Bayesian optimization, achieving convergence on target metrics—such as charge carrier mobility in organic semiconductors—within unprecedented timescales [1]. More recent autonomous laboratories have extended this paradigm to inorganic materials, integrating ab initio databases, literature-mined heuristics, and active learning to navigate vast compositional spaces efficiently [2]. Across these cases, what once required months or years of human trial-and-error can now be achieved in days through iterative, data-informed decisions [3, 13].

Conceptually, this shift reframes the scientific method itself. Linear hypothesis testing gives way to adaptive exploration, in which predictive models—such as graph neural networks for property estimation—are continuously updated using realtime experimental feedback [4, 14]. Human scientists increasingly operate at the level of goal specification and oversight rather than at the level of direct experimental engagement. While this transformation promises efficiency, it simultaneously intensifies ethical stakes by embedding values, assumptions, and trade-offs into autonomous decision architectures.

The literature increasingly recognizes that sustainability narratives surrounding acceleration are incomplete without accounting for computational and infrastructural costs. Training large-scale ML models and operating robotic platforms consume significant energy and water resources, while hardware dependencies introduce vulnerabilities linked to critical material supply chains [6, 15, 16]. Similarly, labor automation concentrates expertise and authority, raising concerns about workforce stratification and long-term innovation resilience [7, 17]. Dual-use risks and governance challenges intersect with these dynamics, particularly when accelerated discovery outpaces ethical deliberation and regulatory adaptation [8, 18-20].

Crucially, these ethical axes do not operate independently. Sustainability justifications may be mobilized to rationalize labor reductions; dual-use safeguards may impose disproportionate burdens on under-resourced institutions; governance gaps may exacerbate inequities in access and influence. Existing ethical frameworks rarely account for such coupled dynamics or for the role of temporal compression in amplifying unintended consequences [21-25].

By articulating acceleration as a multiplier rather than a marginal modifier, this work advances a conceptual synthesis tailored to AI-driven materials science. It provides a structured logic for anticipating and navigating ethical trade-offs before they crystallize into entrenched practices. In doing so, it contributes to ongoing debates on responsible innovation by foregrounding the material, temporal, and epistemic specificities of accelerated discovery. Ultimately, the value of materials acceleration depends not only on how fast materials can be discovered, but on whether such discovery proceeds with ethical foresight sufficient to sustain scientific legitimacy and societal trust [26-29].

Sustainability claims in accelerated materials discovery

Sustainability narratives frequently position materials acceleration as intrinsically resource-efficient, emphasizing reduced experimental iterations, algorithmically guided search, and the rapid identification of environmentally beneficial materials such as next-generation perovskites, battery cathodes, and catalysts for clean energy systems [5, 15]. Within this framing, acceleration is presented as a means to minimize waste, reduce energy expenditure per discovery, and align materials science with broader sustainability imperatives.

However, closer conceptual scrutiny reveals persistent tensions between local experimental efficiencies and system-level resource demands. The operation of self-driving laboratories depends on energy-intensive machine learning model training, continuous computational inference, robotic actuation, and high-throughput characterization infrastructure, all of which carry substantial energy, water, and hardware footprints [6, 16]. These costs are often externalized or omitted from sustainability accounting, creating a risk that gains achieved at the level of experimental iteration are offset—or even surpassed—by upstream and downstream resource consumption.

The emerging literature on “green AI” underscores this concern, arguing that claims of sustainability in AI-enabled systems are incomplete without transparent reporting of computational energy use, carbon intensity, and hardware lifecycles [17]. When applied to accelerated materials discovery, this critique takes on added significance: acceleration increases not only the efficiency of individual discovery cycles, but also their frequency. As discovery timelines compress, the rate at which new materials are proposed, optimized, and deployed correspondingly increases, potentially amplifying cumulative environmental impacts rather than mitigating them.

These dynamics are further complicated by feedback loops between discovery speed and material deployment. Faster identification of promising materials can accelerate scaling and commercialization before full lifecycle impacts—such as mining of critical precursors, supply chain vulnerabilities, or end-of-life disposal—are adequately understood [18]. Conceptually, acceleration thus risks shifting sustainability assessment from a precautionary, lifecycle-oriented practice to a reactive, post-deployment exercise. This tension highlights the need to distinguish between procedural efficiency and systemic sustainability, a distinction often blurred in prevailing acceleration narratives.

Labor displacement and deskilling in automated R&D

Automation in materials research is frequently justified as a means of alleviating repetitive, time-intensive tasks, allowing human researchers to focus on higher-level design, interpretation, and validation. In practice, SDLs redistribute labor away from hands-on synthesis and characterization toward algorithm supervision, platform design, and exception handling [7, 19]. While this redistribution may enhance productivity, it also introduces risks of broader deskilling within the scientific workforce.

Peer-reviewed analyses frame this transformation through the lens of complementarity versus substitution. In principle, AI systems can augment human capabilities by handling routine optimization while preserving human oversight and creativity. In practice, however, expertise may become increasingly concentrated among a narrow cohort of algorithm designers and platform architects. At the same time, technicians, experimentalists, and early-career researchers experience reduced opportunities for skill development [20]. Over time, this can erode tacit knowledge—experiential understanding of material behavior, failure modes, and experimental nuance—that has historically played a critical role in scientific insight and innovation.

These labor effects are not evenly distributed. Displacement and deskilling disproportionately affect junior researchers and technical staff, who are least empowered to influence platform design and most vulnerable to automation without parallel reskilling pathways [21]. Conceptually, this raises ethical concerns not only about employment but about the long-term resilience of the scientific knowledge base. A workforce optimized for supervising autonomous systems may lack the embodied expertise required to detect anomalies, challenge model assumptions, or innovate outside predefined search spaces.

Dual-Use risks in advanced materials technologies

Many materials targeted by accelerated discovery pipelines possess inherent dual-use characteristics. Advances in energy storage materials can simultaneously support renewable energy integration and enable compact, high-density power sources for military or surveillance applications; similarly, high-performance alloys and functional materials developed for civilian infrastructure may enhance defense systems or weapons platforms [8, 22].

Acceleration intensifies these dual-use risks by rapidly generating and disseminating knowledge of high-value materials. Automated pipelines, open datasets, and transferable optimization strategies reduce friction in knowledge transfer, complicating traditional mechanisms of ethical review, export control, and application-specific oversight [23]. At the same time, conventional research timelines allowed for incremental assessment and contextual deliberation, compressed discovery cycles shorten the window for ethical reflection and governance intervention.

Conceptually, the challenge lies not in identifying dual-use potential—long recognized in materials science—but in managing its amplification under conditions of speed and scale. Acceleration transforms dual-use from a downstream concern into an upstream design problem, implicating choices about objective functions, performance metrics, and openness at the earliest stages of discovery. Without anticipatory governance, sustainability-oriented acceleration may inadvertently lower barriers to misuse.

Inequity and governance challenges

Materials acceleration infrastructures are capital-intensive, requiring sophisticated robotics, advanced computation, curated datasets, and interdisciplinary expertise. As a result, access to SDLs and MAPs is concentrated within well-resourced institutions, exacerbating existing global and institutional inequities in research capacity [10]. This concentration risks entrenching discovery asymmetries, whereby a small number of actors control the pace, direction, and benefits of accelerated innovation.

Governance frameworks have not kept pace with these developments. Despite growing recognition of ethical risks, there remain few domain-specific protocols addressing the distinctive challenges posed by accelerated materials R&D, particularly those arising from compressed timelines, autonomous decision-making, and physical experimentation [9]. Calls for inclusive and adaptive governance emphasize the need to involve diverse stakeholders—across regions, disciplines, and career stages—in setting objectives, defining acceptable trade-offs, and distributing benefits [26].

However, inequity itself undermines governance capacity. Under-resourced actors are often excluded from decision-making processes and lack the institutional leverage to shape norms, standards, or oversight mechanisms. This creates a reinforcing cycle in which acceleration concentrates power, weakens collective oversight, and further marginalizes those most affected by its consequences.

Integrative synthesis

Taken together, these ethical axes reveal a tightly coupled system of tensions. Sustainability justifications may obscure labor costs; labor displacement can exacerbate inequity; dual-use risks intensify demands on governance structures that are themselves unevenly distributed [28]. Treating these issues in isolation risks partial or counterproductive interventions. This synthesis underscores the need for an integrated conceptual framework capable of capturing interdependencies, feedback dynamics, and trade-offs, thereby setting the stage for the Ethical Acceleration Tension Matrix introduced in the following section.

Proposed conceptual framework: The Ethical Acceleration Tension Matrix (EATM)

The Ethical Acceleration Tension Matrix (EATM) is proposed as a novel conceptual framework for analyzing and navigating the ethical interdependencies introduced by accelerated materials discovery. In contrast to prevailing AI ethics checklists, principle-based guidelines, or static risk matrices, EATM is explicitly designed to capture dynamic amplification effects arising from compressed timelines and autonomous decision loops in materials acceleration platforms and self-driving laboratories. Rather than treating ethical concerns as discrete or additive, the framework conceptualizes acceleration itself as a structuring force that reshapes how ethical tensions emerge, interact, and propagate.

At the core of EATM is the recognition that acceleration is not ethically neutral. Acceleration (A) functions as a scalar multiplier that intensifies pressures across five analytically distinct but interdependent axes: Sustainability (S), Labor (L), Dual-Use (D), Inequity (I), and Governance (G). Each axis represents a class of ethical concern recurrent in the literature, yet insufficiently examined in combination. Under accelerated conditions, interventions targeting one axis often generate second-order effects on others, producing feedback dynamics that cannot be adequately understood through linear or siloed analysis.

Core logic and theoretical principles

The foundational logic of EATM rests on three interrelated principles.

(1) Tension multiplicity

Ethical axes in accelerated materials discovery do not operate independently. Instead, they form irreducible clusters in which tensions are co-produced. For example, sustainability-driven automation intended to reduce material waste may simultaneously intensify labor displacement by reducing demand for experimental technicians, while also increasing dual-use risk through rapid dissemination of high-performance materials. Isolating a single axis—such as sustainability—without accounting for its interactions systematically distorts ethical evaluation.

(2) Feedback velocity

Acceleration compresses not only experimental timelines but also ethical response windows. As discovery cycles shorten, the pace at which tensions accumulate and propagate increases, reducing opportunities for reactive mitigation. Feedback velocity thus becomes a defining feature of accelerated systems: negative externalities can escalate faster than traditional governance, labor adaptation, or sustainability assessment mechanisms can respond. This necessitates preemptive ethical reasoning embedded in platform design, objective formulation, and access control, rather than post hoc correction.

(3) Equilibrium steering

EATM reframes ethical intervention as a problem of steering complex systems toward more stable and socially desirable equilibria, rather than eliminating trade-offs outright. Ethical governance is conceptualized as the identification and manipulation of leverage points—such as inclusive objective-setting, transparent benchmarking, or mandatory reskilling pathways—that can redirect acceleration dynamics toward positive-sum outcomes. Crucially, equilibrium steering recognizes that zero-risk or zero-conflict states are unattainable under acceleration; the goal is managed balance rather than ethical optimization in isolation.

Structural organization of the matrix

Formally, EATM is organized as a 5×5 interaction matrix, with each axis represented both as a row (primary effects) and a column (secondary amplifications). Diagonal elements capture first-order impacts of acceleration on each ethical axis, while off-diagonal elements encode pairwise interactions and antagonisms. For instance, the Sustainability–Labor (S–L) interaction captures how efficiency-driven automation can deskill experimental labor, while also creating conditional opportunities for reskilling through AI-assisted training systems. Higher-order interactions—where three or more axes reinforce one another—are conceptually represented through the matrix's aggregation around a central node.

This central node, termed the acceleration core, represents compressed timelines and automated decision loops as the generative source of ethical amplification. The core does not merely influence the axes; it is shaped by them, reflecting how governance choices, labor structures, and levels of inequity modulate the magnitude and direction of acceleration's effects.

Tetrahedral representation and non-planar interdependence

To capture the non-linear and non-planar nature of these interdependencies, EATM is further expressed through a tetrahedral geometry **(Figure 1)**. In this representation, Sustainability, Labor, Dual-Use, and Inequity form the vertices of the tetrahedral base, while governance occupies the apex, reflecting its integrative and steering function. Acceleration resides at the volumetric center, symbolizing its role as an omnidirectional amplifier rather than a single-axis driver.

The tetrahedral structure conveys two critical insights. First, no single ethical axis can dominate without destabilizing the system as a whole; perturbations at any vertex propagate volumetrically across the structure. Second, governance is not external to acceleration but structurally embedded within it, shaping how tensions circulate and whether feedback loops become virtuous or vicious. This geometry thus visualizes ethical interdependence as a spatial, rather than hierarchical, problem.

Conceptual contribution

By modeling acceleration as an ethical multiplier with feedback velocity and volumetric interdependence, EATM advances beyond static ethical frameworks toward a dynamic, systems-oriented logic tailored to AI-driven materials science. The framework does not prescribe specific policies or technical solutions. Instead, it provides a generative scaffold for anticipating ethical tensions, comparing intervention strategies, and reasoning about trade-offs under conditions of speed and autonomy. In doing so, EATM offers a conceptual tool for researchers, platform designers, and policymakers seeking

https://doi.org//w306302585

to align accelerated materials discovery with sustainable, equitable, and socially responsible outcomes. **Figure 1** illustrates the tetrahedral structure of the EATM and its internal feedback dynamics.

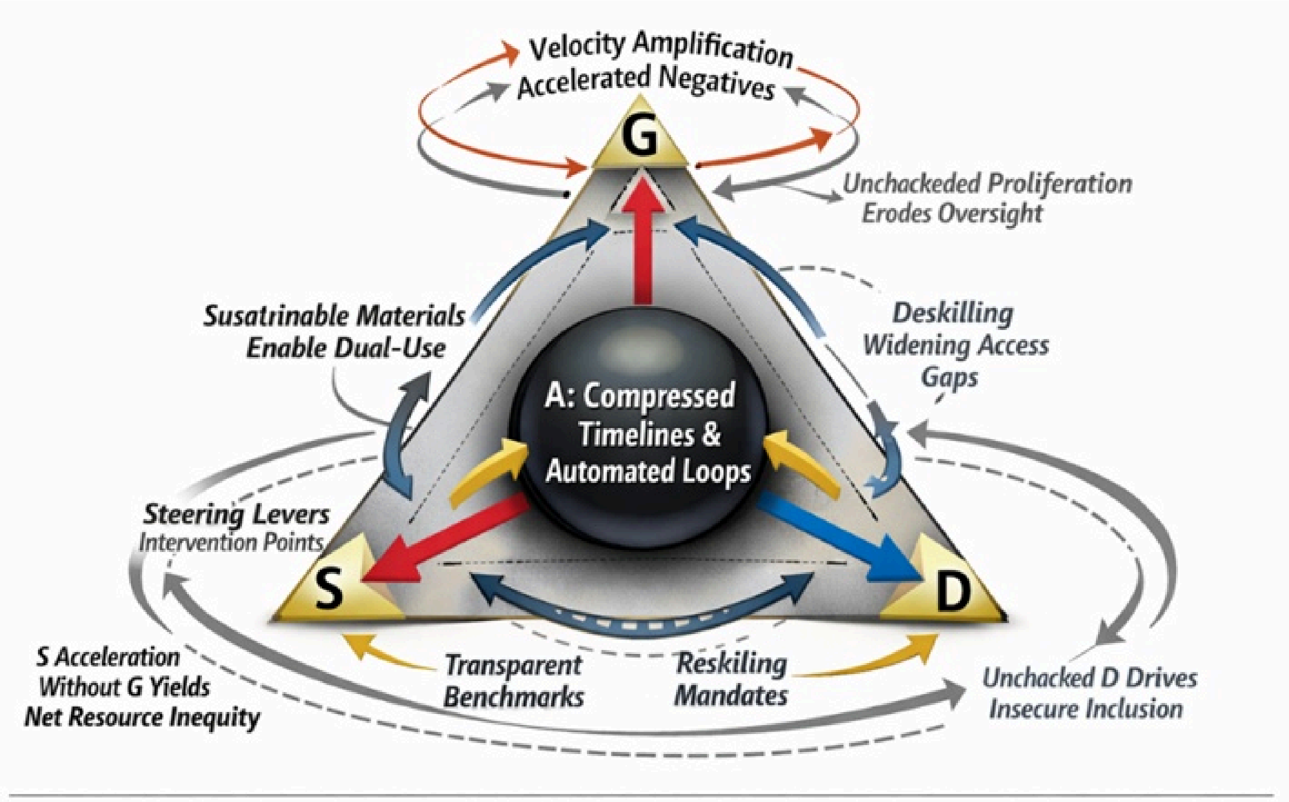

Tetrahedral representation of ethical interdependencies in materials acceleration. *Arrows denote directional influences and feedback; core sphere embodies acceleration dynamics.*

Figure 1. Ethical interdependencies in materials acceleration are illustrated as a tetrahedral model

This framework advances originality by embedding temporal compression and material-physical constraints into ethical logic, offering a tool for researchers, policymakers, and designers to simulate conceptual scenarios (e.g., prioritizing G interventions to dampen S-D synergies). It avoids prescriptive rules and instead provides a generative scaffold for context-specific application.

Analytical implications of the ethical acceleration tension matrix

The ethical acceleration tension matrix (EATM) enables a systematic re-interpretation of ethical challenges in accelerated materials discovery by reframing them as dynamic interactions rather than isolated risks. Through its multidimensional structure, the framework yields several analytical implications that clarify how compressed timelines and autonomous decision loops simultaneously reshape sustainability, labor, dual-use, inequity, and governance.

First, the framework reveals that sustainability claims in accelerated materials discovery are intrinsically conditional rather than absolute. Apparent gains in efficiency—such as reduced experimental iterations or optimized material selection—translate into genuine sustainability benefits only when governance mechanisms explicitly constrain computational intensity, infrastructure scaling, and lifecycle externalities. In the absence of such constraints, acceleration tends to reallocate resource consumption rather than reduce it, producing feedback loops in which speed-driven discovery amplifies cumulative environmental burdens. Sustainability, within EATM, is therefore not a property of acceleration itself but an outcome contingent on governance-mediated boundary conditions.

Second, EATM clarifies how labor transformation under automation extends beyond displacement to encompass deeper epistemic reconfiguration. As SDLs prioritize algorithmic optimization and closed-loop autonomy, experimental labor is progressively abstracted away from direct material engagement. Without structurally embedded reskilling pathways, this shift concentrates epistemic authority among a limited set of platform designers and algorithm specialists, while eroding distributed, hands-on expertise. The framework highlights that labor outcomes are inseparable from inequity dynamics: when reskilling and participation are unevenly accessible, acceleration reinforces hierarchical knowledge structures rather than democratizing innovation.

Third, the framework elucidates the non-linear propagation of dual-use risk under acceleration. When sustainability-oriented objectives align with high-performance material specifications—such as energy density, durability, or efficiency—the same accelerated pathways that enable beneficial applications can simultaneously lower barriers to military or surveillance uses. EATM shows that dual-use risk does not scale linearly with discovery output; instead, it intensifies through interaction with sustainability narratives and openness norms, particularly when governance mechanisms lag behind acceleration velocity. Dual-use thus emerges as a structural property of accelerated systems, not a downstream anomaly.

Fourth, inequity appears in EATM not merely as unequal access to infrastructure, but as a compounding force that weakens ethical oversight itself. High capital requirements for SDL deployment concentrate discovery capacity within elite institutions, while under-resourced actors are excluded from both participation and governance. This asymmetry undermines collective accountability: those most affected by accelerated innovation often have the least influence over its direction, objectives, and safeguards. The framework, therefore, positions inequity as both an ethical axis and a modifier of all other axes, amplifying sustainability, labor, and dual-use tensions when left unaddressed.

Finally, EATM reframes governance as an active steering function rather than a reactive constraint. Traditional regulatory approaches, which operate on extended temporal cycles, are poorly matched to the compressed feedback dynamics of accelerated discovery. The framework highlights the necessity of governance mechanisms that operate at comparable temporal resolution to acceleration itself—such as anticipatory objective-setting, embedded auditability, and interdisciplinary foresight integrated directly into platform design. Ethical governance, in this view, must be temporally synchronized with acceleration rather than imposed retrospectively. Key interaction dynamics among ethical axes and associated leverage points for ethical steering under acceleration are synthesized in Table 2.

Table 2. Inter-axis interactions and ethical steering opportunities in accelerated materials discovery

Interacting axes	Amplification mechanism under acceleration	Risk if unmitigated	Illustrative ethical steering lever
S–L	Efficiency-driven automation reduces experimental labor	Deskilling and loss of tacit scientific knowledge	Mandatory reskilling pathways embedded in SDL platform design
S–D	High-performance sustainability targets align with military utility	Rapid proliferation of dual-use materials	Early-stage objective vetting and controlled dissemination protocols
L–I	Automation favors algorithm designers over experimentalists	Concentration of epistemic authority	Inclusive training access and distributed platform governance
D–G	Open data and rapid optimization outpace oversight	Weak export control and misuse prevention	Anticipatory governance integrated into optimization objectives
I–G	Resource asymmetry limits participation in ethical oversight	Marginalized actors are excluded from norm-setting	Shared infrastructure models and international governance forums
S–G	Sustainability claims outpace	Net increase in resource	Computational and infrastructure

	lifecycle accountability	consumption	footprint auditing

Taken together, these analytical implications underscore the central insight of the EATM: acceleration functions as an ethical multiplier. It intensifies trade-offs, shortens mitigation windows, and transforms localized decisions into system-wide consequences. Addressing ethical challenges in accelerated materials discovery, therefore, requires interventions that acknowledge interdependence, feedback velocity, and equilibrium steering, rather than isolated optimization of individual ethical dimensions. This perspective provides the conceptual foundation for evaluating and guiding responsible acceleration in applied artificial intelligence for materials science.

Results and Discussion

The ethical acceleration tension matrix (EATM) provides a lens for interpreting materials acceleration not as a neutral increase in efficiency, but as a structural transformation that reshapes ethical dynamics through speed, scale, and autonomy. While self-driving laboratories and materials acceleration platforms promise substantial gains in discovery throughput, the analysis demonstrates that these gains are inseparable from intensified ethical tensions best understood through EATM's multidimensional logic. In contrast to domain-agnostic AI ethics frameworks that emphasize abstract principles such as transparency or fairness in isolation, EATM foregrounds the material specificity of accelerated discovery—synthesis safety, physical scalability, computational infrastructure, and resource materiality—which collectively heighten feedback velocity and ethical interdependence [11].

From this perspective, sustainability emerges as a conditional outcome rather than an inherent property of acceleration. Although SDLs/MAPs are frequently framed as reducing waste through optimized experimentation, their operation relies on energy-intensive machine learning workflows, robotic orchestration, and hardware infrastructures dependent on critical minerals [6, 16]. When these costs are not explicitly governed, acceleration reallocates environmental burden rather than reducing it. EATM highlights governance as the integrative axis through which sustainability claims must be audited, ensuring that computational and infrastructural demands are evaluated alongside material performance gains.

Labor dynamics further illustrate how acceleration reconfigures epistemic practice. As autonomous systems assume responsibility for routine experimentation, human roles shift toward oversight, validation, and system design. While this redistribution can enhance efficiency, it also risks eroding tacit, hands-on knowledge central to materials intuition and creative hypothesis formation [7, 17]. EATM clarifies that labor impacts cannot be assessed independently of inequity: without inclusive reskilling pathways and access to platform-level participation, automation concentrates epistemic authority within a narrow subset of actors. This concentration risks innovation bottlenecks, particularly in under-resourced contexts where opportunities for retraining and system-level influence are limited [30, 31].

Dual-use risk is similarly intensified under accelerated conditions. Materials optimized for sustainability—such as high-performance battery electrolytes or photonic materials—may simultaneously enable military or surveillance applications, especially when optimization objectives prioritize performance metrics that are application-agnostic [8, 19]. EATM conceptualizes this risk as volumetric propagation: perturbations at the sustainability–dual-use interface cascade across the system unless governance mechanisms intervene at early design stages. In accelerated, open-data ecosystems, traditional downstream oversight is insufficient; ethical control must be embedded within objective selection, access policies, and dissemination practices [30].

Inequity operates both as a distinct ethical axis and as a force that weakens collective oversight. Capital-intensive infrastructure requirements concentrate accelerated discovery capacity within elite institutions, potentially entrenching dependencies on imported advanced materials and widening global discovery gaps [10, 24]. EATM reveals that inequity amplifies other tensions by excluding marginalized actors from governance processes, thereby reducing the diversity of perspectives shaping acceptable trade-offs. Conversely, interventions at governance–inequity intersections—such as open protocols, shared infrastructure models, and capacity-building mandates—can transform acceleration from a concentrator of power into a potential democratizer of discovery [26].

Across these domains, the central insight of EATM is that isolated interventions generate distortions. Sustainability-driven automation without labor safeguards can exacerbate deskilling; openness without dual-use oversight can accelerate misuse; governance imposed without attention to inequity can reinforce asymmetries [28]. Ethical challenges in accelerated materials discovery, therefore, cannot be resolved through single-axis optimization. Instead, they require coordinated steering across interdependent axes, attuned to the shortened response windows imposed by acceleration itself [9, 25, 31].

The limitations of this analysis are inherent to its conceptual nature. The framework does not provide empirical validation or quantitative predictions, and future work will be needed to test EATM dynamics through case studies, comparative platform analyses, or policy evaluations. Nonetheless, by embedding temporal compression and material constraints directly into ethical reasoning, EATM advances a form of systems-oriented ethics that extends beyond reformulations of general AI guidelines [13, 32-34]. Its value lies in providing a structured scaffold for anticipating trade-offs, designing interventions, and guiding responsible innovation in applied AI for materials science.

Conclusion

This manuscript has introduced the Ethical Acceleration Tension Matrix as a novel conceptual framework for examining ethical interdependencies in accelerated materials discovery. Synthesizing peer-reviewed literature, this study shows how compressed timelines and autonomous decision loops intensify sustainability tensions, labor reconfiguration, dual-use vulnerability, inequity, and governance challenges through irreducible feedback dynamics. By representing these interactions within a tetrahedral logic, the framework captures ethical complexity that is obscured by linear or checklist-based approaches.

The analysis underscores that the benefits commonly attributed to materials acceleration are contingent on proactive ethical integration rather than post hoc mitigation. Without deliberate steering, efficiency gains risk being undermined by unaccounted computational burdens, erosion of experimental expertise, proliferation of misuse potential, and widening access disparities. Responsible acceleration, therefore, requires governance mechanisms capable of operating at the same temporal resolution as discovery itself, alongside labor and equity interventions that preserve epistemic diversity and societal trust.

Beyond its immediate application to self-driving laboratories and materials acceleration platforms, EATM offers a transferable logic for other domains of AI-enabled scientific discovery where physical experimentation, high-stakes outcomes, and rapid iteration converge. Future conceptual extensions may incorporate evolving autonomy levels, hybrid human–AI epistemologies, or comparative governance regimes. At the same time, practical translation will demand domain-specific protocols attuned to the material realities of experimentation and deployment.

Ultimately, the framework advances a central claim: acceleration is an ethical multiplier. Its value for materials science lies not only in how quickly discoveries are made, but in how responsibly speed, autonomy, and societal impact are aligned. By providing a structured conceptual foundation for this alignment, the ethical acceleration tension matrix helps ensure that accelerated materials discovery serves sustainable, equitable, and socially legitimate ends in applied artificial intelligence.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 30 Dec 2023 Revised: 24 Mar 2024 Accepted: 16 Apr 2024
Published online: 18 July 2024

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

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.

Burger B, Maffettone PM, Gusev VV, Aitchison CM, Bai Y, Wang X, et al. A mobile robotic chemist. Nature. 2020;583(7815):237-41.

Coley CW, Eyke NS, Jensen KF. Autonomous discovery in the chemical sciences: Part I—Progress. Angew Chem Int Ed. 2020;59(51):22858-93.

Tabor DP, Roch LM, Saikin SK, Kreisbeck C, Sheberla D, Montoya JH, et al. Accelerating the discovery of materials for clean energy in the era of smart automation. Nat Rev Mater. 2020;5(1):5-20.

Gongora AE, Xu B, Perry W, Okoye C, Riley P, Reyes KG, et al. A Bayesian experimental autonomous researcher for mechanical design. Sci Adv. 2020;6(15):eaaz1708.

Roch LM, Häse F, Kreisbeck C, Tamayo-Mendoza T, Yunker LP, Hein JE, et al. ChemOS: Orchestration software to democratize autonomous discovery. PLoS One. 2020;15(4):e0229862.

Häse F, Roch LM, Aspuru-Guzik A. Olympus: A benchmarking framework for noisy optimization and experiment planning. Mach Learn Sci Technol. 2021;2(3):035021.

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

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.

Epps RW, Bowen MS, Volk AA, Abdel-Latif K, Abolhasani M. Artificial intelligence for autonomous reaction systems. Matter. 2022;5(11):3644-60.

Henderson P, Hu J, Romero J, Brunskill E, Jurafsky D, Pineau J. Towards the systematic reporting of the energy and carbon footprints of machine learning. J Mach Learn Res. 2020;21:1-43.

Strubell E, Ganesh A, McCallum A. Energy and policy considerations for deep learning in NLP. Proc ACL. 2020;1:3645-50.

Luccioni AS, Schmidt V, Bengio Y. Estimating the carbon footprint of machine learning. Patterns. 2022;3(6):100493.

Schwartz R, Dodge J, Smith NA, Etzioni O. Green AI. Commun ACM. 2020;63(12):54-63.

Whittaker M, Crawford K, Dobbe R, et al. AI Now Report 2020: The social implications of automation. AI Now Institute. 2020.

Benjamins R, Barbado A, Sierra D. Responsible AI by design in practice. AI Ethics. 2021;1:87-97.

Canty RB, Li R, Salazar R. Benchmarking autonomous laboratories. Nat Commun. 2023;14:6402.

Kleinberg J, Ludwig J, Mullainathan S, Obermeyer Z. Prediction policy problems. Am Econ Rev. 2020;110(5):491-5.

National Academies of Sciences, Engineering, and Medicine. Dual Use Research of Concern in the Life Sciences. Washington DC: National Academies Press; 2020.

Brundage M, Avin S, Clark J, Toner H, EckersleyP, Garfinkel B, et al. The malicious use of artificial intelligence: Forecasting, prevention, and mitigation. AI Magazine. 2020;41(2):1-17.

Kalluri P. Don’t ask if artificial intelligence is good or fair, ask how it shifts power. Nature. 2020;583:169.

Jobin A, Ienca M, Vayena E. The global landscape of AI ethics guidelines. Nat Mach Intell. 2019;1:389-99.

Floridi L, Cowls J, Beltrametti M, Chatila R, Chazerand P, Dignum V, et al. AI4People—An ethical framework for a good AI society. Minds Mach. 2020;28(4):689-707.

Stilgoe J, Owen R, Macnaghten P. Developing a framework for responsible innovation. Res Policy. 2013;42(9):1568-80.

Delgado-Licona F, Abolhasani M. Research acceleration in self-driving laboratories. Adv Intell Syst. 2023;5(3):2200331.

Abolhasani M, Kumacheva E. The rise of self-driving laboratories in chemical and materials sciences. Nat Synth. 2023;2:483-92.

Rai A, Constantinides P, Sarker S. Next-generation digital platforms. MIS Q. 2019;43(1):3-9.

Beck S, Mahony M. The politics of anticipation. Environ Polit. 2018;27(1):1-20.

Jasanoff S. A field of its own: The emergence of science and technology studies. STS Handbook. 2020 edition.

Häse F, Roch LM, Aspuru-Guzik A. Next-generation experimentation with self-driving laboratories. Trends Chem. 2019;1(3):282–91.

Stach E, DeCost B, Kusne AG, Hattrick-Simpers J, Brown KA, Reyes KG, et al. Autonomous experimentation systems for materials development: A community perspective. MRS Bull. 2021;46(8):656–63.

Liu R, Gadepally V, Krishnamurthy A, et al. The data bottleneck: Machine learning in scientific discovery. Patterns. 2021;2(5):100256.

Vinuesa R, Azizpour H, Leite I, Balaam M, Dignum V, Domisch S, et al. The role of artificial intelligence in achieving the Sustainable Development Goals. Nat Commun. 2020;11:233.

Saltelli A, Bammer G, Bruno I, Charters E, Di Fiore M, Didier E, et al. Five ways to ensure that models serve society: A manifesto. Nature. 2020;582:482–4.