ORIGINAL RESEARCH Open access

Delegated Experimentation: Decision Authority Transfer in Closed-Loop Materials Engineering

Natalia Petrova[1*], Elena Stoyanova[1], Ivan Dimitrov[2]

Abstract

Closed-loop systems have become foundational to computational and data-driven materials engineering, integrating automated experimentation, machine learning inference, and orchestration software to compress the design-make-test-analyze cycle. These pipelines rely on continuous flows of data, models, and decisions, yet the mechanisms governing the transfer of decision authority between human experts and autonomous agents remain conceptually underdeveloped. Existing infrastructures emphasize optimization and execution but offer limited interpretive frameworks for how authority is dynamically delegated across epistemic states and pipeline stages. This manuscript presents the Decision Authority Delegation Cascade (DADC) Framework, an original systems-level architecture that formalizes delegated experimentation as a structured cascade of authority transfer. The framework delineates layered pipelines—from data representation through model inference and steering logics to execution and feedback—while emphasizing infrastructure trade-offs in representation fidelity, uncertainty quantification, and delegation thresholds. Synthesizing advances in Bayesian active learning, self-driving laboratories, and orchestration platforms, the DADC Framework interprets authority transfer not as a binary handover but as a continuous, computationally steered process that modulates discovery dynamics. The framework offers interpretive insights into scalable computational ecosystems, highlighting pathways to align human epistemic oversight with autonomous operation and to mitigate bottlenecks in closed-loop materials discovery. Its application reframes infrastructure design around explicit delegation logics, with implications for the next generation of autonomous materials platforms.

Keywords Bayesian optimization, Self-driving laboratories, Delegated experimentation, Decision authority transfer, Closed-loop materials discovery, Computational orchestration

*Correspondence:
Natalia Petrova
natalia.petrova@outlook.com
[1] Department of Computational Materials Engineering, Faculty of Engineering, University of Sofia, Sofia, Bulgaria
[2] Department of Data-Driven Materials Systems, Faculty of Engineering, Technical University of Sofia, Sofia, Bulgaria

Introduction

The emergence of closed-loop systems in computational materials engineering

Computational materials science has shifted from static simulation-driven screening toward integrated, iterative discovery pipelines. Closed-loop architectures couple high-throughput experimentation, real-time characterization, and machine-learning models to enable autonomous iteration. These systems treat materials discovery as a feedback-controlled process in which each cycle refines both the candidate space and the underlying models. Early demonstrations established the feasibility of on-the-fly optimization in complex synthesis–structure–property landscapes, showing that Bayesian strategies can guide experimentation without exhaustive enumeration [1-3]. Subsequent platforms extended this paradigm to thin-film optimization, superconducting materials, and organic synthesis, demonstrating that closed-loop operation can function at scales unattainable by manual workflows [4-6].

The infrastructure underpinning these systems has matured rapidly. Orchestration software now coordinates robotic hardware, analytical instruments, and computational agents, while data infrastructures support streaming representation of experimental outcomes into model-update routines. This convergence has produced ecosystems in which the boundary between computation and physical experimentation is increasingly porous [7, 8].

Data-driven paradigms and the autonomy spectrum

Data-driven materials engineering operates along a spectrum of autonomy. At one end, human experts define experimental spaces and interpret results; at the other, fully autonomous platforms propose, execute, and evaluate campaigns with minimal intervention. Between these poles lies a continuum of hybrid operation in which decision authority is shared. Self-driving laboratories exemplify this hybridity, employing modular robotic platforms, online analytics, and planning algorithms to execute multi-step workflows [9-11]. Mobile robotic chemists and flow-based synthesis robots have further illustrated how physical execution can be decoupled from human presence while remaining responsive to computational guidance [12, 13].

Despite these advances, the literature reveals a persistent conceptual focus on optimization algorithms and hardware integration rather than on the governance of decision rights. Most platforms treat authority allocation as an implicit design choice rather than an explicit, tunable system property [14, 15]. Consequently, current infrastructures lack standardized logics for when and how authority should shift from human oversight to automated steering. This gap limits scalability: as pipelines grow in complexity, the cognitive load of constant human supervision becomes unsustainable, while premature delegation risks epistemic misalignment [11, 16].

Infrastructure challenges in authority allocation

Closed-loop systems generate three interrelated challenges for decision authority. First, data representation must capture not only measured properties but also epistemic uncertainty and contextual metadata. Second, inference engines must propagate uncertainty into actionable confidence metrics that can inform delegation. Third, steering logics must balance exploration, exploitation, and safety constraints under varying levels of delegated authority. These challenges are compounded by the heterogeneous nature of materials data—sparse, noisy, and multimodal—which complicates reliable uncertainty quantification [1, 3, 17].

Current orchestration approaches address parts of this triad but rarely integrate them into a unified authority-transfer model. Robotic platforms demonstrate impressive execution autonomy, yet human operators typically retain veto power or define campaign boundaries [4, 9]. Bayesian optimizers excel at suggesting candidates yet operate within pre-defined search spaces set by humans [2, 17]. The result is a patchwork of authority regimes rather than a coherent cascade.

Positioning the conceptual contribution

This manuscript addresses the above gap by introducing the Decision Authority Delegation Cascade (DADC) Framework as an original interpretive architecture for delegated experimentation. The framework treats decision authority transfer as the central organizing principle of closed-loop materials engineering, structuring the discovery pipeline into interdependent layers that explicitly modulate authority at each transition. It provides a systems-level lens for analyzing representation–inference interactions, feedback dynamics, and infrastructure trade-offs without invoking empirical benchmarks or testable hypotheses. The DADC Framework thereby reframes autonomous materials platforms as computational governance systems in which authority delegation becomes a steerable design parameter.

Theoretical Background & Literature Synthesis

Closed-loop experimentation and active learning

Closed-loop materials discovery has emerged as a defining paradigm within computational and data-driven materials engineering, crystallizing from the convergence of active learning, surrogate modeling, and automated experimentation [1, 3]. At its core, the closed-loop logic reframes discovery not as a static screening exercise but as a dynamically adaptive inference process. Active learning algorithms—most prominently Bayesian optimization frameworks—enable efficient traversal of vast compositional and process spaces by explicitly modeling uncertainty alongside predicted performance [2, 17]. Acquisition functions such as expected improvement, probability of improvement, and information gain operationalize the exploration–exploitation trade-off, directing experimental resources toward regions that maximize epistemic value rather than purely predicted optima [2, 18].

On-the-fly implementations have demonstrated that surrogate models—ranging from Gaussian processes to deep kernel learners and neural operators—can be iteratively retrained as new experimental data streams are acquired [1, 3]. This continuous updating transforms the experimental campaign into a sequential decision system in which each measurement reshapes the topology of the search landscape. Rather than predefining exhaustive sampling grids, closed-loop pipelines dynamically reallocate effort, concentrating resources in high-uncertainty or high-promise regimes. This epistemic steering capacity is particularly consequential in high-dimensional materials systems where combinatorial explosion renders brute-force exploration computationally and experimentally infeasible [1, 5].

Subsequent research extended Bayesian active learning beyond property optimization into domains such as catalytic reaction tuning, crystal-structure assembly, phase diagram reconstruction, and defect engineering [6, 19, 20]. In these contexts, uncertainty-aware selection consistently demonstrated accelerated convergence relative to heuristic or random sampling. Importantly, these systems do not merely predict outcomes; they actively shape the experimental trajectory. Model confidence, posterior variance, and predictive entropy become operational signals that directly influence laboratory action [2, 3]. The discovery pipeline thus becomes an inference engine embedded in physical reality—one where data acquisition and model evolution co-produce knowledge through recursive interaction.

Self-Driving laboratories and robotic orchestration

If active learning provides the cognitive architecture of closed-loop discovery, self-driving laboratories constitute its physical instantiation. These platforms integrate synthesis, processing, and characterization capabilities within modular, software-orchestrated infrastructures [9-11]. Robotic deposition systems fabricate thin-film libraries; automated flow reactors conduct reaction optimization; robotic arms assemble metal–organic frameworks and polymer networks [4, 19]. Instrumentation pipelines feed characterization outputs—spectroscopy, diffraction, imaging—directly into computational analysis layers, closing the data loop between execution and inference [4, 5].

Recent advances have expanded these systems beyond fixed automation cells into mobile and reconfigurable robotic ecosystems. Autonomous mobile robots can navigate laboratory environments, transporting samples between instruments, reconfiguring workflows, and decoupling experimental logic from static spatial layouts [12]. This mobility transforms the laboratory itself into a programmable substrate, where physical routing becomes an adjustable parameter within the discovery algorithm.

Central to this orchestration layer are campaign management platforms such as ChemOS and related autonomy software stacks [7, 8]. These systems standardize communication between planning algorithms, robotic hardware, and analytical instrumentation through unified application programming interfaces. By abstracting hardware heterogeneity, orchestration software enables researchers to define high-level experimental objectives—target properties, compositional constraints, optimization goals—without scripting individual machine actions [7]. The platform translates these objectives into executable task sequences, effectively compiling scientific intent into robotic instruction sets.

Despite these advances, execution autonomy remains bounded. Human researchers continue to define objective functions, permissible chemical spaces, safety envelopes, and termination criteria [9, 11]. These boundary conditions function as governance scaffolds, delimiting the operational sovereignty of the autonomous system. Thus, while laboratories increasingly execute experiments independently, the authority to define what constitutes a meaningful experiment remains partially centralized in human decision structures.

Hybrid human–autonomous interfaces

Parallel to the expansion of physical autonomy is the evolution of hybrid decision architectures that embed human expertise within algorithmic loops. Human-in-the-loop active learning systems allow domain experts to inject qualitative knowledge—feasibility judgments, mechanistic priors, categorical exclusions—into Bayesian optimization pipelines [21]. Such interventions can reshape acquisition landscapes, preventing the system from pursuing chemically implausible or operationally unsafe regions despite favorable statistical signals.

Emerging work further explores symbolic–subsymbolic integration. Reactivity-seeking neural networks, knowledge-graph-guided optimizers, and large-language-model (LLM) agents have been deployed to generate hypotheses, propose synthesis routes, and interpret anomalous outcomes [22-24]. In these configurations, symbolic reasoning augments numerical optimization, enabling the system to operate across both quantitative and semantic knowledge layers. For instance, LLM agents can translate literature insights into experimental constraints, while graph reasoning engines infer plausible reaction pathways that guide robotic execution [23].

Case studies in autonomous chemical experimentation repeatedly underscore the enduring necessity of human oversight [25, 26]. Edge cases—unexpected phase transitions, equipment anomalies, hazardous reaction pathways—require contextual judgment that exceeds current algorithmic generalization. Moreover, conceptual innovation—redefining target properties, reframing discovery questions, or identifying paradigm-shifting anomalies—remains deeply human-driven. Hybrid systems therefore function not as replacements but as collaborative cognition platforms, distributing reasoning tasks across biological and artificial agents.

Yet, while the literature richly documents these hybrid interactions, it rarely formalizes the mechanics of authority transfer. Human guidance is described descriptively —“expert input,” “manual override,” “constraint injection”— without computational representation. Delegation remains implicit, embedded in interface design rather than modeled as a quantifiable system variable [11, 14].

Synthesis: Current state of computational ecosystems

Taken collectively, the contemporary closed-loop ecosystem demonstrates remarkable maturity across three infrastructural axes: data acquisition, model inference, and physical execution. High-throughput experimentation pipelines generate continuous data streams; machine learning architectures transform these streams into predictive knowledge; robotic laboratories enact algorithmically guided experimentation at accelerating speeds [1, 9, 15]. Discovery has thus become an integrated cyber-physical process in which sensing, reasoning, and acting operate within tightly coupled feedback loops.

However, this infrastructural sophistication reveals an asymmetry. While throughput optimization, uncertainty quantification, and novelty detection have received extensive formalization, the governance of decision authority has not. Who—or what—decides the next experiment? Under what conditions is authority escalated, shared, or revoked? How are epistemic risk, safety constraints, and strategic priorities encoded into delegation structures?

At present, these decisions are embedded heuristically within software settings, campaign definitions, or human supervisory practices [14, 16]. Authority allocation is treated as an operational convenience rather than a computational primitive. This absence of formal delegation architectures limits interpretability, scalability, and trust—particularly as laboratories progress toward higher degrees of autonomy.

This synthesis foregrounds a critical infrastructure opportunity: the integration of explicit authority governance into the computational backbone of closed-loop materials discovery. Embedding delegation logics alongside data, models, and robotics would enable traceable decision rights, adaptive control hierarchies, and quantifiable autonomy gradients.

The Delegated Autonomous Discovery Control (DADC) Framework emerges in response to this gap. By conceptualizing authority transfer as a layered, modelable process—spanning objective definition, experimental prioritization, execution authorization, and exception handling—DADC reframes autonomy not as a binary state but as a dynamically negotiated continuum. In doing so, it provides both an analytical lens for interrogating existing discovery systems and a design architecture for constructing future computational ecosystems in which intelligence and authority co-evolve.

Proposed conceptual framework

The Decision Authority Delegation Cascade (DADC) Framework conceptualizes delegated experimentation as a multi-layered pipeline in which decision authority transfers progressively from human-centric to machine-centric control. The framework comprises four structural layers connected by data–model–discovery pipelines and bidirectional feedback loops.

The Acquisition and Representation Layer ingests raw experimental outputs and contextual metadata, producing structured representations that encode both measured values and epistemic provenance. These representations feed the Inference and Uncertainty Layer, where surrogate models generate predictions together with spatially and temporally resolved uncertainty estimates. The Delegation and Steering Layer evaluates these epistemic signals against predefined but tunable criteria to compute a delegation fraction—the proportion of decision authority assigned to autonomous agents for the next cycle. Finally, the Execution and Feedback Layer translates delegated decisions into robotic actions or campaign adjustments, closing the loop by returning new observations to the acquisition layer.

Data → Model → Discovery pipelines operate continuously: experimental data update surrogate models, which inform steering logics, which in turn modulate execution parameters. Feedback loops exist at two scales. Local loops within a single campaign refine model parameters in near real time; global loops across campaigns adjust delegation thresholds and representation schemas. Computational steering logics are implemented at the delegation layer through continuous comparison of system confidence against human-assigned reliability benchmarks. When model uncertainty falls below a dynamic threshold, authority shifts toward automation; when unexpected outcomes increase epistemic divergence, authority reverts toward human oversight. The DADC Framework can be visualized as a vertical cascade **(Figure 1)**.

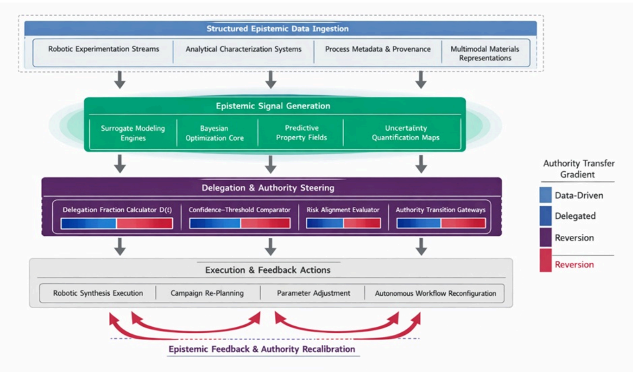

Figure 1. Decision Authority Delegation Cascade (DADC): Layered Architecture of Authority Transfer in Closed-Loop Materials Discovery

It architecture illustrating the progressive transfer of decision authority across closed-loop materials discovery pipelines. Structured experimental data enter through the Acquisition and Representation Layer and propagate into inference systems that generate predictive and uncertainty signals. These epistemic metrics feed the Delegation and Steering Layer, where authority fractions are dynamically computed using confidence-threshold comparisons. Delegated decisions are executed through robotic platforms, while bidirectional feedback loops recalibrate representation fidelity, inference reliability, and governance thresholds. The cascade emphasizes authority transfer as a continuous computational variable embedded within discovery infrastructure rather than as a binary human–machine handover.

The dynamics of authority transfer within the DADC Framework can be conceptualized as

$$D(\tau) = \frac{1}{1 + \exp\left(-\frac{C\_sys(\tau) - \theta(\tau)}{\sigma}\right)} \quad (1)$$

where D(τ) is the delegated authority fraction at pipeline step τ, Csys(τ) is the composite system confidence derived from model predictions and uncertainty estimates, θ(τ) is a time- or context-dependent delegation threshold set by higher-level campaign objectives, and σ is a sensitivity parameter controlling the sharpness of the transition. This logistic form captures the nonlinear, threshold-driven nature of authority transfer: small changes in epistemic alignment near the threshold produce rapid shifts in autonomy, while regions far from the threshold remain stable.

The interaction between delegated authority and feedback-loop strength may be expressed as

$$\frac{dK}{dt} = \alpha \cdot G(\tau) \cdot D(\tau) + \beta \cdot (1 - D(\tau)) \quad (2)$$

where K represents the cumulative knowledge state of the pipeline, G(τ) is the instantaneous information gain from new observations, α and β are scaling coefficients reflecting the relative efficiency of autonomous versus human-guided learning, and D(τ) modulates the contribution of each mode. This differential relation interprets feedback as authority-weighted: greater delegation amplifies the impact of machine-driven information gain, while retained human authority preserves a baseline learning rate grounded in expert intuition.

Finally, the epistemic alignment between representation fidelity and inference confidence can be captured by the interaction term

$$E = \int \left(Rf\left(x\right) - Ic\left(x\right)\right) \cdot w\left(D\right)\, dx \quad (3)$$

where R_f is the fidelity of the data representation, I_c is the inference confidence, and w(D) is a weighting function that increases with delegated authority. The integral quantifies cumulative misalignment across the candidate space; steering logics act to minimize E by adjusting either representation schemas or delegation levels. Each layer performs a distinct computational and governance function within the cascade **(Table 1).**

Table 1. Structural Layers of the Decision Authority Delegation Cascade (DADC) Framework

Layer	Primary Function	Computational Components	Autho
Acquisition & Representation	Structured data ingestion and epistemic encoding	Experimental streams, metadata capture, multimodal fusion	No aut epi sub fori
Inference & Uncertainty	Predictive modeling and confidence estimation	Surrogate models, Bayesian optimizers, uncertainty fields	In au influe con m
Delegation & Steering	Authority allocation and governance computation	Threshold comparators, delegation fraction estimators	Core tra comp
Execution & Feedback	Physical experimentation and campaign adaptation	Robotics, orchestration software, workflow engines	Opera del au

These relations are not prescriptive equations but interpretive devices that formalize the computational steering logics embedded in the DADC Framework. They illustrate how authority transfer functions as both a consequence of and a driver for pipeline dynamics, enabling infrastructure designers to treat delegation as an explicit, optimizable variable within closed-loop materials engineering.

Analytical implications

The Decision Authority Delegation Cascade (DADC) Framework generates a set of interconnected analytical implications that reorient the design logic of closed-loop materials engineering infrastructures. By elevating decision authority transfer to a first-class computational primitive, the framework converts previously implicit design choices into explicit, steerable parameters distributed across the four layers.

Representation–inference alignment and epistemic risk structures

The Acquisition and Representation Layer determines the upper bound on inference reliability. Any loss in representational fidelity propagates directly into the uncertainty fields of the Inference and Uncertainty Layer, increasing the integral misalignment term E. Within the DADC lens, elevated E functions as an automatic trigger for authority reversion: the Delegation and Steering Layer responds by lowering D(τ), thereby retaining human oversight until representation schemas are refined or additional metadata streams are incorporated. This mechanism interprets epistemic risk not as a static uncertainty budget but as a dynamic, pipeline-wide state variable that governs authority allocation in real time.

The implication is architectural: future data infrastructures must embed provenance tracking and multimodal fusion routines at the representation stage, because any downstream delegation decision is only as robust as the fidelity of the incoming data structures. Platforms that treat representation as a preprocessing step decoupled from delegation logic will encounter hard ceilings on safe autonomy [1, 15].

Feedback dynamics and authority-weighted learning

The differential relation

$$\frac{dK}{dt} = \alpha \cdot G(\tau) \cdot D(\tau) + \beta \cdot (1 - D(\tau)) \quad (4)$$

formalizes the differential efficiency of autonomous versus human-guided information uptake. When D(τ) is high, the pipeline capitalizes on the rapid, parallel sampling capacity of robotic execution; when D(τ) is low, it preserves the slower but conceptually richer integration characteristic of expert intuition. The analytical consequence is that optimal long-term knowledge accumulation trajectories are unlikely to be monotonic in autonomy. Instead, the framework predicts oscillatory delegation patterns—periods of high delegation for broad exploration followed by targeted human re-engagement when epistemic divergence exceeds threshold θ(τ).

This oscillation insight carries direct infrastructure consequences. Orchestration layers must support reversible delegation handovers without loss of campaign state, including the ability to replay prior decision contexts when authority reverts. Current orchestration software that assumes unidirectional autonomy ramps will require extension with bidirectional state synchronization primitives [7, 8].

Steering logics and infrastructure trade-offs

The logistic Delegation Function

$$D(\tau) = \frac{1}{1 + \exp\left(-\frac{C\_sys(\tau) - \theta(\tau)}{\sigma}\right)} \quad (5)$$

reveals that small adjustments to the threshold θ(τ) \theta(\tau) θ(τ) or sensitivity σ \sigma σ produce nonlinear shifts in pipeline behavior. Raising θ(τ) \theta(\tau) θ(τ) (more conservative delegation) increases robustness to model misspecification and safety-critical edge cases but reduces experimental throughput. Lowering θ(τ) \theta(\tau) θ(τ) accelerates discovery campaigns yet amplifies the risk of cascading errors when surrogate models encounter out-of-distribution regimes. The framework therefore frames infrastructure design as a multi-objective optimization over the delegation surface rather than over material performance alone.

A further implication concerns modularity. Because each layer exposes well-defined interfaces (representation schemas, uncertainty fields, delegation fractions, execution primitives), the cascade can be decomposed and recomposed across institutions or hardware platforms. A thin-film deposition group can adopt only the Delegation and Steering Layer while retaining its existing robotic stack; a high-throughput synthesis consortium can share the Inference and Uncertainty Layer across distributed laboratories. The DADC architecture thereby supports federated discovery ecosystems in which authority transfer becomes the interoperability standard [15, 16]. These governance relations formalize delegation as a computationally tractable infrastructure variable **(Table 2)**.

Table 2. Mathematical Governance Functions in the DADC Authority-Transfer Model

Relation	Functional Role	Governing Variables	Infrastru Interpre
Logistic Delegation Function D(τ)	Computes delegated authority fraction	System confidence, threshold θ, sensitivity σ	Autho transi regul
Knowledge Accumulation Differential dK/dt	Models authority-weighted learning	Information gain G, delegation fraction D	Learr efficie modul
Representation–Inference Integral E	Quantifies epistemic misalignment	Representation fidelity Rf, inference confidence Ic	Risk det metı
Threshold Sensitivity Parameter σ	Controls delegation volatility	Confidence–threshold gradient	Govern stability

Scalability and long-term discovery governance

At campaign scales exceeding thousands of experiments, the cognitive burden of constant human supervision becomes prohibitive. The DADC Framework interprets this scalability barrier as a delegation-capacity limit: the human operator functions as a finite-bandwidth governor whose intervention rate must be budgeted across the cascade. By making delegation explicit and quantifiable, the framework enables infrastructure designers to allocate that finite bandwidth optimally—reserving human attention for high-impact epistemic pivots while delegating routine optimization to autonomous agents [9, 11].

Collectively, these implications reposition closed-loop platforms from optimization engines to governance systems. The central analytical contribution of DADC is to replace ad-hoc authority allocation with a layered, mathematically interpretable cascade that renders every major design decision—representation format, uncertainty propagation, threshold tuning, feedback topology—visible and adjustable.

The Decision Authority Delegation Cascade (DADC) Framework sits at the intersection of three maturing infrastructural threads in computational materials engineering: Bayesian active learning pipelines, robotic orchestration platforms, and hybrid human–machine decision systems. Each of these domains has independently advanced the autonomy frontier of materials discovery, yet their integration has remained largely operational rather than governance-centered. Bayesian optimizers have transformed experimental prioritization through uncertainty-aware acquisition logics [1, 2, 17], robotic laboratories have automated synthesis and characterization execution at unprecedented throughput [4, 9, 10], and hybrid cognitive architectures have embedded human reasoning within algorithmic discovery loops [21, 23]. The DADC Framework does not seek to extend any single algorithmic component; instead, it introduces the missing authority-governance layer that systemically integrates these infrastructures into a coherent authority-flow architecture.

Existing self-driving laboratory implementations demonstrate impressive execution autonomy yet retain human-defined campaign boundaries, safety envelopes, and termination criteria [9, 11]. Platforms coordinating thin-film optimization, reaction discovery, and superconducting materials exploration routinely rely on human specification of search spaces and objective functions despite automated experimental iteration [4-6]. The DADC cascade renders these boundary conditions computationally explicit, reframing them as tunable governance parameters rather than fixed supervisory constraints. Authority thresholds, delegation fractions, and epistemic divergence triggers become dynamically adjustable system variables embedded within the discovery loop.

Similarly, Bayesian optimization frameworks excel at candidate selection yet typically operate within externally imposed authority regimes. Acquisition functions—expected improvement, entropy search, or categorical optimizers such as Gryffin—prioritize experiments without determining who authorizes execution [2, 17, 18]. By embedding delegation computation directly within the inference layer, DADC allows predictive confidence and uncertainty gradients to modulate autonomy continuously rather than episodically. This integration transforms inference outputs into governance signals rather than purely predictive metrics.

The literature on hybrid intelligence further underscores the persistent need for expert oversight in edge cases. Human-in-the-loop optimization studies demonstrate that domain experts can reshape acquisition landscapes through mechanistic priors and feasibility constraints [21]. Large-language-model agents and symbolic reasoning systems have similarly augmented autonomous discovery by translating literature knowledge into experimental design constraints or mechanistic hypotheses [23, 24]. Autonomous reactivity discovery platforms and robotic chemists repeatedly encounter anomalous regimes—unexpected reaction pathways, metastable phase transitions, equipment anomalies—that demand contextual interpretation beyond current machine generalization [22, 25]. Within the DADC architecture, such epistemic divergences trigger authority reversion through measurable confidence–threshold misalignment, converting qualitative supervisory intervention into a computable control mechanism.

Importantly, the framework is complementary to existing orchestration infrastructures rather than competitive with them. Campaign management platforms such as ChemOS and related autonomy stacks already standardize communication across planning algorithms, robotics, and analytical instrumentation [7, 8]. DADC can be retrofitted into these environments by inserting a Delegation and Steering Layer between inference and execution modules, enabling authority computation without altering underlying hardware or optimization routines. This modular insertion preserves legacy infrastructure investment while extending governance sophistication.

A natural analytical extension concerns multi-agent and multi-institutional cascades. Distributed laboratory networks—spanning synthesis facilities, characterization hubs, and computational modeling centers—are increasingly coordinated through shared orchestration environments [15, 16]. Within such federated ecosystems, delegation fractions **D(τ)** can be computed at the consortium level, enabling coordinated authority handovers across geographic and institutional boundaries. Autonomous mobile robotic chemists and modular automation platforms further enable physically distributed execution layers capable of responding to centrally computed governance signals [12, 13]. The same mathematical structures governing single-pipeline delegation thus scale to multi-node discovery networks in which authority transfer becomes an interoperability protocol.

This federated extension carries implications for infrastructure resilience and knowledge equity. Low-cost or “frugal twin” self-driving laboratory models expand access to autonomous experimentation capabilities across resource-constrained institutions [14]. Embedding delegation governance within such distributed architectures ensures that authority allocation remains transparent and traceable even when experimentation is geographically dispersed or hardware-heterogeneous.

The framework’s limitations are conceptual rather than empirical. The logistic delegation relation and authority-weighted learning differential are interpretive abstractions selected for analytical tractability. Alternative functional representations—including reinforcement-learned delegation policies, piecewise governance rules, or probabilistic voting systems—remain compatible provided they preserve the layered cascade topology. Likewise, the framework deliberately abstracts hardware execution as a black-box primitive. Robotic synthesis platforms, flow reactors, mobile chemists, and crystallization robots differ substantially in throughput, latency, and safety envelopes [4, 12, 26], yet DADC operates at the governance layer above these execution heterogeneities.

This abstraction facilitates cross-platform generality but necessitates domain-specific instantiation when deployed in concrete infrastructures. For example, autonomous crystal assembly systems or robotic MOF synthesis platforms may require tighter delegation thresholds due to sensitivity to environmental perturbations [19, 27], whereas continuous-flow reaction optimizers can tolerate higher autonomy due to stable operating regimes [6, 20].

Ultimately, the primary contribution of the DADC Framework remains infrastructural rather than algorithmic. It supplies a shared conceptual vocabulary and computational backbone for discussing, designing, and comparing authority regimes across the rapidly diversifying landscape of closed-loop materials platforms. By formalizing authority transfer as a quantifiable, model-integrated variable, DADC enables autonomous discovery systems to evolve from execution-optimized pipelines into governance-aware cyber-physical ecosystems in which intelligence, responsibility, and decision rights co-evolve.

Conclusion

The Decision Authority Delegation Cascade (DADC) Framework introduces a new systems-level architecture for delegated experimentation in closed-loop materials engineering. By structuring the discovery pipeline into four interdependent layers connected by explicit authority-transfer mechanisms, the framework converts decision governance from an implicit design feature into a computationally steerable primitive.

Three original interpretive relations—the logistic delegation function, the authority-weighted knowledge accumulation equation, and the representation–inference misalignment integral—formalize the dynamics of authority transfer, feedback amplification, and epistemic alignment. These relations, together with the layered cascade topology, enable infrastructure designers to analyze and optimize closed-loop systems along the dimension of decision rights rather than solely along material performance or throughput.

The analytical implications span representation fidelity requirements, feedback-loop topologies, delegation-threshold tuning, and scalability to federated ecosystems. Collectively, they reposition autonomous materials platforms as computational governance systems in which human epistemic oversight and machine efficiency are balanced through continuous, data-driven authority modulation.

By providing a unified interpretive lens for representation–inference interactions, steering logics, and infrastructure trade-offs, the DADC Framework addresses a foundational gap in the computational materials engineering literature. It supplies the conceptual and architectural scaffolding required to move from isolated demonstrations of autonomy toward scalable, interoperable, and epistemically aligned discovery infrastructures.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 21 Jul 2024 Revised: 15 Aug 2024 Accepted: 16 Nov 2024
Published online: 18 March 2025

References

Kusne AG, Yu H, Wu C, Zhang H, Hattrick-Simpers J, DeCost B, et al. On-the-fly closed-loop materials discovery via Bayesian active learning. Nat Commun. 2020;11:5966. https://doi.org/10.1038/s41467-020-19597-w.

Shields BJ, Stevens J, Li J, Parasram M, Damani F, Alvarado JIM, et al. Bayesian reaction optimization as a tool for chemical synthesis. Nature. 2021;590(7844):89-96.

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.

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

Pogue EA, New A, McElroy K, Le NQ, Pekala MJ, McCue I, et al. Closed-loop superconducting materials discovery. npj Comput Mater. 2023;9:181. https://doi.org/10.1038/s41524-023-01131-3.

Nambiar AMK, Breen CP, Hart T, Kulesza T, Jamison TF, Jensen KF. Bayesian optimization of computer-proposed multistep synthetic routes on an automated robotic flow platform. ACS Cent Sci. 2022;8(6):825-36.

Sim M, Vakili MG, Strieth-Kalthoff F, Hao H, Hickman RJ, Miret S, et al. ChemOS 2.0: An orchestration architecture for chemical self-driving laboratories. Matter. 2024. https://doi.org/10.1016/j.matt.2024.04.022.

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

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

Abolhasani M, Kumacheva E. The rise of self-driving labs in chemical and materials sciences. Nat Synth. 2023;2(5):483-92.

Seifrid M, Pollice R, Aguilar-Granda A, Morgan Chan Z, Hotta K, Ser CT, et al. Autonomous chemical experiments: Challenges and perspectives on establishing a self-driving lab. Acc Chem Res. 2022;55(17):2454-66.

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

Steiner S, Wolf J, Glatzel S, Andreou A, Granda JM, Keenan G, et al. Organic synthesis in a modular robotic system driven by a chemical programming language. Science. 2019;363(6428):eaav2211.

Lo S, Baird SG, Schrier J, Blaiszik B, Carson N, Foster I, et al. Review of low-cost self-driving laboratories in chemistry and materials science: The “frugal twin” concept. Digit Discov. 2024;3(6):842-68.

Stach E, DeCost B, Kusne AG, Hattrick-Simpers J, Brown KA, Reyes KG, et al. Autonomous experimentation systems for materials development: A community perspective. Matter. 2021;4(9):2702-26. https://doi.org/10.1016/j.matt.2021.06.036.

Rahmanian F, Flowers J, Guevarra D, Richter M, Fichtner M, Donnely P, et al. Enabling modular autonomous feedback-loops in materials science through hierarchical experimental laboratory automation and orchestration. Adv Mater Interfaces. 2022;9(21):2101987.

Häse F, Aldeghi M, Hickman RJ, Roch LM, Aspuru-Guzik A. GRYFFIN: An algorithm for Bayesian optimization of categorical variables informed by expert knowledge. Appl Phys Rev. 2021;8(3):031406.

Aldeghi M, Häse F, Hickman RJ, Tamblyn I, Aspuru-Guzik A. Golem: An algorithm for robust experiment and process optimization. Chem Sci. 2021;12(44):14792-807.

Xie Y, Zhang C, Deng H, Zheng B, Su J-W, Shutt K, et al. Accelerate synthesis of metal-organic frameworks by a robotic platform and Bayesian optimization. ACS Appl Mater Interfaces. 2021;13(45):53485-91.

Clayton AD, Pyzer-Knapp EO, Purdie M, Jones MF, Barthelme A, Pavey J, et al. Bayesian self-optimization for telescoped continuous flow synthesis. Angew Chem Int Ed. 2023;62(28):e202214511.

Adams F, McDannald A, Takeuchi I, Kusne AG. Human-in-the-loop for Bayesian autonomous materials phase mapping. Matter. 2024:7(2):697-709.

Caramelli D, Granda JM, Mehr SHM, Cambié D, Henson AB, Cronin L. Discovering new chemistry with an autonomous robotic platform driven by a reactivity-seeking neural network. ACS Cent Sci. 2021;7(11):1821-30.

Boiko DA, MacKnight R, Kline B, Gomes G. Autonomous chemical research with large language models. Nature. 2023;624(7992):570-8.

Zhu Q, Zhang F, Huang Y, Xiao H, Zhao L, Zhang X, et al. An all-round AI-chemist with a scientific mind. Natl Sci Rev. 2022;9(12):nwac190.

Granda JM, Donina L, Dragone V, Long D-L, Cronin L. Controlling an organic synthesis robot with machine learning to search for new reactivity. Nature. 2018;559(7714):377-81.

Duros V, Grizou J, Xuan W, Hosni Z, Long D-L, Miras HN, et al. Human versus robots in the discovery and crystallization of gigantic polyoxometalates. Angew Chem Int Ed. 2017;56(36):10815-20.

Masubuchi S, Morimoto M, Morikawa S, Onodera M, Asakawa Y, Watanabe K, et al. Autonomous robotic searching and assembly of two-dimensional crystals to build van der Waals superlattices. Nat Commun. 2018;9:1413.