

ORIGINAL RESEARCH

Open access

Autonomous Laboratories and Closed-Loop Materials Discovery: A Conceptual Framework for Self-Driving Experimentation

Thabo Nkosi^{1*}, Lerato Molefe¹, Siphso Dlamini², Ayanda Mokoena¹, Kabelo Ndlovu²

Abstract

Materials discovery remains one of the central bottlenecks in the translation of scientific ideas into functional technologies. Conventional workflows often depend on slow cycles of hypothesis, synthesis, testing, interpretation, and redesign. These cycles are constrained by human time, manual dexterity, instrument availability, and the practical impossibility of exploring vast compositional and processing spaces exhaustively. Incremental improvements in automation have accelerated particular laboratory tasks, yet they have not fully transformed the logic of discovery. A robot that performs a predefined protocol faster than a human still operates within a human-directed trial-and-error paradigm. The deeper opportunity lies in laboratories that can decide what to do next, execute experiments, learn from outcomes, and adapt their strategy without waiting for complete human redesign. This article proposes a conceptual framework for autonomous, closed-loop laboratories in materials discovery. The framework treats discovery as an adaptive, self-improving process rather than a linear sequence of isolated experiments. It argues that autonomous laboratories are not merely high-throughput platforms but emerging scientific systems in which artificial intelligence, robotics, characterization, data infrastructure, and human judgement are integrated through a shared feedback logic. The framework defines the closed-loop discovery cycle, identifies six conceptual layers of an autonomous laboratory, and explains how experiment selection, robotic execution, real-time data capture, and human oversight interact. It also highlights why data quality, machine-readability, and error handling are foundational rather than peripheral concerns. The resulting vision is of materials discovery as a learning ecosystem in which each experimental outcome improves not only the model of a material system but also the strategy by which future knowledge is generated.

Keywords Conceptual framework, Materials discovery, Active learning, Closed-loop discovery, Autonomous laboratory, Self-driving experimentation

*Correspondence:

Thabo Nkosi

thabo.nkosi@gmail.com

¹ Department of Computational Materials Science, Faculty of Engineering and the Built Environment, University of Cape Town, Cape Town, South Africa

² Department of Data-Driven Materials Systems, Faculty of Engineering, University of the Witwatersrand, Johannesburg, South Africa

Introduction

Materials discovery has traditionally advanced through a human-directed sequence of intuition, formulation, synthesis, characterization, and interpretation. This workflow has produced extraordinary scientific progress, yet Stach and colleagues [1] describe how it remains limited by slow experimental iteration, fragmented data, and the difficulty of coordinating computation, synthesis, and measurement at scale. As design spaces expand from individual compounds to multicomponent compositions, processing histories, interfaces, and device architectures, the number of plausible

experiments rapidly exceeds what any human team can manually evaluate. The result is a discovery process that is not only slow but also unevenly biased toward familiar chemistries, accessible equipment, and hypotheses that fit existing mental models.

The autonomous laboratory emerges as a response to these structural limitations rather than simply as a faster version of the conventional laboratory. Tom and colleagues [2] frame self-driving laboratories as integrated systems that connect decision algorithms, automated experimentation, and learning from data, while Abolhasani and Kumacheva [3] emphasise their significance for chemical and materials sciences as a new mode of research organisation. In this view, autonomy does not mean removing scientists from discovery; it means embedding scientific reasoning into an iterative system that can test, learn, and redirect itself. The laboratory becomes an adaptive platform whose behaviour changes as evidence accumulates.

The conceptual distinction between automation and autonomy is central to this article. Automation executes a known procedure with reduced manual intervention, whereas autonomy uses feedback to determine which procedure should be executed next, a distinction made concrete in discussions of autonomous chemical experimentation by Seifrid and colleagues [4]. A high-throughput robot may generate many samples, but without adaptive selection, model updating, and feedback, it remains a production tool rather than a self-driving discovery system. Autonomous laboratories therefore require an integrated logic that joins artificial intelligence, robotics, characterization, and data systems into a closed loop.

This article articulates a conceptual framework for closed-loop, self-driving materials discovery by synthesising evidence and emerging architectures across the field. It first defines the closed-loop discovery cycle and then proposes a layered architecture for autonomous laboratories, building on community perspectives on materials development [1] and recent orchestration architectures such as ChemOS 2.0 [5]. Subsequent sections examine active learning, robotics, data infrastructure, human oversight, and the challenges that must be overcome for autonomous discovery to mature. The aim is to clarify the underlying paradigm: materials discovery as a self-improving ecosystem rather than a collection of accelerated laboratory tasks.

The Closed-Loop Discovery Cycle

The closed-loop discovery cycle begins when a scientific goal is translated into an operational search problem that a laboratory system can pursue. In autonomous materials platforms, this cycle links hypothesis generation, experiment selection, automated execution, characterization, data capture, model updating, and re-iteration, as illustrated by closed-loop materials discovery through Bayesian active learning [6]. Each loop transforms experimental outcomes into improved beliefs about the material system and improved decisions about the next experiment. The essential feature is not speed alone but continuity between prediction and action.

Closed-loop discovery narrows the gap between computational recommendation and experimental reality. MacLeod and colleagues [7] demonstrated this logic in accelerated thin-film materials discovery, where autonomous experimentation connected decision-making, synthesis, and measurement within a single adaptive workflow. A later self-driving laboratory study advancing the Pareto front for material properties [8] showed how the loop can pursue multiple objectives rather than optimise a single scalar target. Conceptually, this means the laboratory no longer waits for a human to interpret each result before the search trajectory changes.

The closed loop also reframes failure as information. In conventional experimentation, failed syntheses, noisy measurements, and unexpected outcomes are often treated as interruptions, but autonomous workflows can treat them as signals that refine the model, expose constraints, or redirect the search. The “minimal working example” proposed by Baird and Sparks [9] is important because it makes clear that a self-driving laboratory must include decision autonomy, experiment execution, and feedback, not merely a collection of instruments. A closed-loop laboratory is therefore a scientific learning system whose intelligence is expressed through iteration.

Figure 1 illustrates the adaptive closed-loop discovery cycle that connects scientific goals, experiment selection, robotic execution, characterization, data capture, model updating, and renewed experimental decision-making.




Figure 1. Closed-loop materials discovery as an adaptive self-driving experimentation cycle.

Architecture of an Autonomous Laboratory

An autonomous laboratory can be conceptualised as a layered system in which each layer contributes a necessary function to the closed loop. The AI decision layer defines candidate experiments, selects actions, and updates models; the robotic execution layer physically realises those actions; the characterization layer measures outcomes; the data-management layer captures and structures evidence; the feedback and optimization layer converts observations into updated strategy; and the human oversight layer defines goals, constraints, and interpretive boundaries. This layered view is consistent with community-level descriptions of autonomous experimentation systems [1] and with architecture-focused work on orchestration in self-driving laboratories [5]. The value of the framework lies in showing that autonomy emerges from interaction among layers rather than from any single algorithm or instrument.

Figure 2 presents the layered architecture through which AI decision-making, robotic execution, characterization, data management, feedback optimization, and human oversight jointly produce laboratory autonomy.




Figure 2. Layered architecture of an autonomous laboratory for self-driving materials discovery.

The AI decision layer is the strategic centre of the system, but it cannot function in isolation. ChemOS was developed as orchestration software to democratize autonomous discovery [10], and ChemOS 2.0 extends this idea by emphasising coordination among agents, instruments, workflows, and decision modules [5]. These orchestration systems suggest that the architecture of autonomy is partly computational and partly organisational: models must communicate with equipment, data stores, and human supervisors in a form that supports timely action. Without orchestration, even powerful algorithms remain disconnected recommendations rather than laboratory behaviour.

The robotic execution and characterization layers form the physical interface between computational intent and material reality. Burger and colleagues [11] showed how a mobile robotic chemist could perform experimental work in a laboratory environment, while Dave and colleagues [12] connected robotic experimentation and machine learning for autonomous battery electrolyte discovery. These examples indicate that the physical layer must be more than a rigid automation line; it must be sufficiently flexible to translate algorithmic choices into reliable synthesis, preparation, and measurement operations. Characterization, in turn, must produce data quickly and consistently enough for the loop to continue learning.

The data-management, feedback, and human oversight layers bind the architecture into an accountable scientific system. Data management matters because autonomous laboratories require machine-readable records, traceable metadata, and reliable information flow, a point emphasised by Willoughby and Frey [13] and extended in knowledge-graph approaches to laboratory automation [14]. Human oversight remains essential because scientists define meaningful objectives, safety boundaries, and interpretations that cannot be reduced to optimization metrics alone. Table 1 outlines the six conceptual layers of an autonomous laboratory architecture and their functions.

Table 1. Conceptual Architecture of an Autonomous Laboratory: Layers, Functions, and Interactions

Conceptual layer	Primary function	Key interactions within the closed loop	Conceptual contribution to autonomy
AI decision layer	Converts scientific goals into candidate experiments and prioritised actions	Receives structured data from prior experiments and passes selected actions to orchestration and robotics	Provides adaptive direction by deciding what the laboratory should do next
Robotic execution layer	Performs synthesis, formulation, sample handling, preparation, and protocol execution	Receives machine-readable instructions from the decision and orchestration layers and returns execution status	Turns computational intent into reproducible physical action
Characterization layer	Measures properties, structures, performance, and quality indicators	Supplies experimental outcomes to the data-management and feedback layers	Creates the evidence on which model updating and scientific interpretation depend
Data-management layer	Captures raw data, metadata, context, provenance, and error records	Connects instruments, robots, models, databases, and human users through structured information flow	Makes experimental outcomes reusable, traceable, and machine-readable
Feedback and optimization layer	Updates models, evaluates uncertainty, and balances exploration with exploitation	Integrates measurements and metadata to refine the next experimental decision	Closes the loop by converting results into improved strategy

Human oversight layer	Defines goals, constraints, ethics, safety limits, and interpretive priorities	Supervises objectives, reviews anomalies, and adjusts the scope of autonomy	Ensures that autonomous discovery remains scientifically meaningful and responsible
-----------------------	--	---	---

Active Learning and Experiment Selection

Active learning is the decision intelligence that makes closed-loop experimentation more than automated repetition. In a conventional workflow, scientists often select experiments through intuition, factorial design, or incremental variation around known formulations, but Bayesian active learning can select experiments according to expected information gain and uncertainty reduction, as demonstrated in closed-loop materials discovery by Kusne and colleagues [6]. The core conceptual shift is that the next experiment is not chosen because it is merely available or familiar. It is chosen because it is expected to improve the laboratory's understanding of the search space.

Bayesian optimization gives the autonomous laboratory a principled way to balance exploration and exploitation. In autonomous materials synthesis in flow, Epps and colleagues [15] showed how artificial intelligence can guide experimental choices under practical constraints, while robotic metal–organic framework synthesis using Bayesian optimization illustrates the selection of experiments in complex synthetic spaces [16]. Exploration directs the system toward uncertain regions where learning may be high, whereas exploitation directs it toward regions likely to yield improved performance. A self-driving laboratory must continually negotiate between these two impulses.

Active learning also changes what is meant by efficiency in materials discovery. Efficiency is not simply the number of experiments performed per day; it is the amount of useful knowledge gained per experimental action, a principle reflected in hierarchical active learning of nonequilibrium phase diagrams [17]. In multiobjective self-driving experimentation, MacLeod and colleagues [8] showed that discovery can target trade-offs rather than a single optimum. This matters because materials problems often involve competing requirements such as conductivity, stability, processability, cost, and sustainability.

The experiment-selection layer therefore acts as the cognitive engine of the autonomous laboratory. Tom and colleagues [2] describe self-driving laboratories as systems that integrate decision-making with experimental execution, and the minimal working example proposed by Baird and Sparks [9] clarifies that autonomy requires an algorithmic choice about what to do next. The accelerated synthesis of novel materials in an autonomous laboratory further shows how algorithmic decision-making can guide the search toward experimentally realised compounds [18]. Conceptually, active learning transforms discovery from a passive accumulation of results into an adaptive strategy for asking better questions.

Robotics and Automation as the Physical Layer

Robotics provides the physical embodiment of the autonomous laboratory's decisions. A model may identify the next promising experiment, but the loop only becomes real when robotic systems translate that recommendation into synthesis, sample preparation, treatment, transfer, and measurement. Burger and colleagues [11] demonstrated the symbolic importance of a mobile robotic chemist by showing that laboratory action itself can be reorganised around autonomous operation. Similarly, autonomous battery electrolyte discovery linked robotic experimentation with machine learning, making the physical layer part of the learning process rather than a separate service function [12].

The physical layer must be flexible enough to support diverse materials chemistries rather than optimised only for a narrow repetitive task. Low-cost self-driving laboratory concepts, including the “frugal twin” approach, suggest that modularity and accessibility are important if autonomous discovery is to spread beyond highly resourced laboratories [19].

Integrating autonomy into automated research platforms also requires interfaces that allow instruments, software, and workflows to communicate without constant bespoke engineering [20]. Robotic inorganic synthesis guided through complex phase diagrams further illustrates the need for automation that can navigate variable recipes, conditions, and outcomes [21].

Robotics also exposes the materiality of autonomy: powders clog, liquids evaporate, films crack, instruments drift, and samples fail. Bayley and colleagues [22] emphasise that autonomous chemistry must navigate real laboratory messiness rather than assume idealised execution, while autonomous thin-film discovery showed that synthesis and characterization must be reliable enough to sustain iterative learning [7]. Battery electrolyte optimization through robotic experimentation likewise demonstrates that closed-loop systems depend on repeatable execution as much as on intelligent selection [23]. The physical layer is therefore not a passive actuator but a source of constraints, data, and uncertainty that shapes the entire discovery cycle.

Data as the Nervous System

Data is the nervous system of the autonomous laboratory because it carries signals between models, robots, instruments, databases, and human scientists. Willoughby and Frey [13] argue that data management is central to digital discovery, not a secondary administrative task. In an autonomous setting, every experimental action must generate machine-readable evidence that can be interpreted without waiting for retrospective human reconstruction. Without structured data flow, the closed loop becomes a set of disconnected events rather than a learning system.

The autonomous laboratory requires more than raw measurements; it requires context. Knowledge-graph approaches to laboratory automation, such as the platform-to-knowledge-graph evolution described by Bai and colleagues [14], show how experimental entities, procedures, materials, instruments, and outcomes can be linked into computationally usable structures. A dynamic knowledge graph for distributed self-driving laboratories further extends this idea by treating data infrastructure as a shared coordination layer across locations and systems [24]. Such structures allow the laboratory to remember not only what happened but also how, where, under what conditions, and with what degree of reliability.

Metadata and provenance are especially important because autonomous systems learn from patterns that may be distorted by hidden experimental context. The materials experiment knowledge graph proposed by Statt and colleagues [25] shows how materials experiments can be represented in ways that support reuse, interpretation, and computational reasoning. Orchestration platforms also depend on real-time information about instrument state, sample identity, workflow progress, and failures, as reflected in integrated autonomy for automated research platforms [20]. In this sense, metadata is not annotation after the fact but part of the experimental signal itself.

A closed-loop laboratory must also treat errors as data. ChemOS 2.0 emphasises orchestration across experimental agents and workflows [5], and this orchestration becomes scientifically meaningful only when failed actions, interrupted runs, drift, and anomalous measurements are captured in forms that models can interpret. Traditional laboratory notebooks often preserve human narrative, whereas autonomous laboratories require structured records that support both machine action and human auditability. Table 2 compares traditional and autonomous laboratory data management practices.

Table 2. Data Management in Traditional Versus Autonomous Laboratories: Data Capture, Metadata, Error Handling, and Machine-Readability

Data-management dimension	Traditional laboratory practice	Autonomous laboratory practice	Implication for closed-loop discovery

Data capture	Measurements are often recorded after experiments through notebooks, spreadsheets, or instrument exports	Data are captured automatically and continuously from instruments, robots, sensors, and workflow software	Enables rapid model updating and reduces dependence on retrospective reconstruction
Metadata	Context may be incomplete, inconsistent, or dependent on individual researcher habits	Experimental conditions, sample identity, instrument state, operator actions, and timing are recorded in structured formats	Makes experiments interpretable, reusable, and comparable across cycles
Error handling	Failed experiments may be omitted, informally described, or treated as unusable	Failures, warnings, abnormal states, and recovery actions are logged as part of the experimental record	Allows models and humans to learn from failure modes rather than ignore them
Machine-readability	Records are often designed primarily for human reading and later interpretation	Records are designed for computational access, querying, reasoning, and automated decision-making	Allows the laboratory to convert outcomes into immediate next-step decisions
Provenance	Links among sample, protocol, instrument, parameter, and result may be difficult to reconstruct	Provenance is maintained across the full experimental workflow	Supports auditability, reproducibility, and trustworthy autonomy
Real-time availability	Data may become usable only after cleaning, transfer, or manual interpretation	Data streams are available during or immediately after execution	Enables the loop to respond dynamically to emerging evidence

Human Scientists in the Loop

Autonomous laboratories do not eliminate human scientists; they transform the location and character of human agency. Human-in-the-loop Bayesian autonomous materials phase mapping shows that scientists can guide discovery by shaping objectives, evaluating plausibility, and intervening when automated search requires domain judgement [26]. Seifrid and colleagues [4] similarly emphasise that autonomous experimentation raises practical and conceptual challenges that require human supervision. The scientist becomes less a manual executor and more a designer of goals, constraints, and meanings.

In this transformed role, the human scientist defines what counts as a valuable discovery. Community perspectives on autonomous experimentation stress the importance of aligning materials development with meaningful scientific and technological objectives [1], while self-driving laboratory reviews show that algorithms require goals that are formalised yet scientifically grounded [2]. A laboratory can optimise a metric, but humans must decide whether that metric adequately represents performance, novelty, safety, sustainability, or manufacturability. The closed loop therefore depends on human judgement at the level of purpose.

Human scientists are also anomaly investigators and interpreters of unexpected behaviour. Autonomous chemistry must handle surprising outcomes, failed reactions, and ambiguous measurements, as Bayley and colleagues [22] note in discussing the navigation of self-driving laboratories. The rise of self-driving labs in chemical and materials sciences further indicates that human expertise remains necessary for framing systems, validating results, and recognising when models are learning the wrong lesson [3]. Autonomy shifts human attention from repetitive execution toward the interpretation of exceptions.

This shift requires new forms of scientific training and collaboration. Researchers must understand enough about algorithms, robotics, metadata, and workflow orchestration to supervise autonomous systems responsibly, a need implied by low-cost self-driving laboratory frameworks [19] and orchestration architectures such as ChemOS 2.0 [5]. The original ChemOS vision also suggests that democratizing autonomous discovery requires usable interfaces rather than tools accessible only to specialists in automation or machine learning [10]. The future scientist in an autonomous laboratory is therefore a hybrid practitioner: part domain expert, part systems thinker, part data steward, and part ethical supervisor.

Challenges and Future Vision

The autonomous laboratory paradigm faces challenges that are technical, organisational, and epistemic. Stach and colleagues [1] identify the need for community-level coordination in autonomous experimentation systems, while Seifrid and colleagues [4] emphasise obstacles in establishing reliable self-driving laboratories. Integrating autonomy into research platforms also reveals practical difficulties in connecting software, instruments, workflows, and decision-making processes [20]. These challenges matter because a closed loop can amplify weaknesses as quickly as it amplifies learning.

Noisy data, failed syntheses, instrument drift, and model bias are not peripheral problems; they are central threats to trustworthy autonomy. Bayesian active learning can guide discovery efficiently, but Kusne and colleagues [6] also show that closed-loop learning depends on the quality of experimental feedback. Robotic inorganic synthesis across complex phase diagrams highlights how challenging real materials spaces can be, especially when phase boundaries, kinetic effects, and synthesis conditions interact [21]. Autonomous synthesis of novel materials further illustrates that discovery platforms must handle uncertainty in both prediction and experimental realisation [18].

A second challenge is transferability across laboratories, materials classes, and workflows. Frugal and low-cost self-driving laboratory concepts indicate that broader adoption will depend on architectures that are modular, affordable, and adaptable [19]. ChemOS 2.0 points toward orchestration as a route to generality, yet distributed self-driving laboratories require shared data structures and interoperable knowledge representations [5, 24]. The field must therefore move from impressive individual demonstrations toward reusable infrastructures that allow autonomous discovery strategies to travel.

The future vision is a self-improving discovery ecosystem that learns not only which materials to make but also which strategies are most effective for particular scientific goals. Hierarchical active learning of nonequilibrium phase diagrams suggests that autonomous systems can learn structured representations of complex materials spaces [17], while accelerated artificial intelligence development for materials synthesis in flow shows how models and experiments can co-evolve [15]. In this vision, the laboratory becomes an adaptive scientific partner whose performance improves through experience, oversight, and accumulated data. Table 3 summarises the key challenges in closed-loop materials discovery and potential mitigation strategies.

Table 3. Challenges in Closed-Loop Autonomous Discovery: Sources, Consequences, and Mitigation Approaches

Challenge	Primary source	Consequence for autonomous discovery	Potential mitigation approach
Noisy or sparse data	Measurement uncertainty, limited experiments, heterogeneous instruments	Models may overfit, underexplore, or select misleading next experiments	Use uncertainty-aware models, replicate critical measurements, and encode measurement quality in metadata
Failed experiments	Unstable synthesis routes, robotic errors, incompatible	The loop may stall or learn from incomplete records	Treat failures as structured data and design recovery protocols

	reagents, sample loss		within orchestration software
Instrument drift	Calibration changes, sensor ageing, environmental fluctuations	Apparent material trends may reflect equipment variation rather than real behaviour	Include calibration logs, drift detection, and routine reference measurements
Model bias	Training data imbalance, narrow search history, inappropriate objective functions	Search may reinforce familiar regions and miss unconventional materials	Combine exploration incentives, human review, and periodic reassessment of objectives
Safety constraints	Reactive chemistries, high temperatures, toxic substances, pressure, electrical hazards	Autonomous action may create unacceptable operational risks	Encode safety boundaries, interlocks, human approval points, and constrained optimization
Lack of standard interfaces	Proprietary instruments, inconsistent data formats, bespoke laboratory software	Workflows become difficult to reproduce, scale, or transfer	Develop modular APIs, shared schemas, and interoperable orchestration layers
Workflow transferability	Differences in materials class, equipment, protocol, and local expertise	Autonomous strategies may work in one laboratory but fail elsewhere	Build adaptable workflows with explicit assumptions, metadata-rich provenance, and validation stages
Human trust and accountability	Opaque algorithms, unclear responsibility, insufficient interpretability	Scientists may resist adoption or accept recommendations uncritically	Maintain human oversight, interpretable decision records, and transparent audit trails

Conclusion

This article has proposed a conceptual framework for understanding autonomous laboratories as closed-loop, learning-driven discovery systems. The framework defines materials discovery not as a linear sequence of manual choices but as an adaptive cycle in which experiment selection, execution, characterization, data capture, model updating, and human oversight reinforce one another. Its central contribution is to organise existing developments in artificial intelligence, robotics, automation, and data infrastructure into a unified paradigm for self-driving experimentation.

The autonomous laboratory should therefore be understood as more than a productivity tool. Its deeper significance lies in the reconfiguration of scientific practice: experiments become decisions made under uncertainty, data become real-time signals for action, robots become physical agents in the learning loop, and scientists become strategic supervisors of adaptive systems. This reconfiguration has the potential to transform materials engineering from slow trial-and-error into a self-improving ecosystem.

Realising this vision will require sustained interdisciplinary collaboration among materials scientists, chemists, roboticists, data engineers, computer scientists, instrument developers, and ethicists. The field must build platforms that are not only fast but also reliable, interpretable, safe, transferable, and scientifically meaningful. If these challenges are met, autonomous laboratories can become a foundational infrastructure for the next era of materials discovery.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 07 May 2024

Revised: 25 Aug 2024

Accepted: 10 Nov 2024

Published online: 18 January 2025

Rights and permissions

Open Access The author(s) retain copyright. This article is licensed under the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License. It may be shared and adapted for non-commercial purposes with appropriate attribution, an indication of changes, and distribution of adaptations under the same license. Third-party material may be subject to separate terms identified in its credit line. View the license at <https://creativecommons.org/licenses/by-nc-sa/4.0/>.

References

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>.

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(6):483-92.

<https://doi.org/10.1038/s44160-022-00231-0>.

Seifrid M, Pollice R, Aguilar-Granda A, Chan ZM, 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.

<https://doi.org/10.1021/acs.accounts.2c00220>.

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;7(9):2959-77.

<https://doi.org/10.1016/j.matt.2024.04.022>.

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(1):5966.

<https://doi.org/10.1038/s41467-020-19597-w>.

MacLeod BP, Parlange 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).

<https://doi.org/10.1126/sciadv.aaz8867>.

MacLeod BP, Parlange FGL, Rupnow CC, Dettelbach KE, Elliott MS, Morrissey TD, et al. A self-driving laboratory advances the Pareto front for material properties. *Nat Commun.* 2022;13(1):995.

<https://doi.org/10.1038/s41467-022-28580-6>.

Baird SG, Sparks TD. What is a minimal working example for a self-driving laboratory? *Matter.* 2022;5(12):4170-8.

<https://doi.org/10.1016/j.matt.2022.11.007>.

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).

<https://doi.org/10.1371/journal.pone.0229862>.

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

<https://doi.org/10.1038/s41586-020-2442-2>.

Dave A, Mitchell J, Kandasamy K, Wang H, Burke S, Paria B, et al. Autonomous discovery of battery electrolytes with robotic experimentation and machine learning. *Cell Rep Phys Sci.* 2020;1(12):100264.

<https://doi.org/10.1016/j.xcrp.2020.100264>.

Willoughby C, Frey JG. Data management matters. *Digit Discov.* 2022;1(3):183-94.

<https://doi.org/10.1039/D1DD00046B>.

Bai J, Cao L, Mosbach S, Akroyd J, Lapkin AA, Kraft M. From platform to knowledge graph: Evolution of laboratory automation. *JACS Au.* 2022;2(2):292-309.

<https://doi.org/10.1021/jacsau.1c00438>.

Epps RW, Volk AA, Reyes KG, Abolhasani M. Accelerated AI development for autonomous materials synthesis in flow. *Chem Sci.* 2021;12(17):6025-36.

<https://doi.org/10.1039/D0SC06463G>.

Xie Y, Zhang C, Deng H, Zheng B, Su JW, 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.

<https://doi.org/10.1021/acsami.1c16506>.

Ament S, Amsler M, Sutherland DR, Chang MC, Guevarra D, Connolly AB, et al. Autonomous materials synthesis via hierarchical active learning of nonequilibrium phase diagrams. *Sci Adv.* 2021;7(51).

<https://doi.org/10.1126/sciadv.abg4930>.

Szymanski NJ, Rendy B, Fei Y, Kumar RE, He T, Milsted D, et al. An autonomous laboratory for the accelerated synthesis of novel materials. *Nature.* 2023;624(7990):86-91.

<https://doi.org/10.1038/s41586-023-06734-w>.

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(5):842-68.

<https://doi.org/10.1039/D3DD00223C>.

Canty RB, Koscher BA, McDonald MA, Jensen KF. Integrating autonomy into automated research platforms. *Digit Discov.* 2023;2(5):1259-68.

<https://doi.org/10.1039/D3DD00135K>.

Chen J, Cross SR, Miara LJ, Cho JJ, Wang Y, Sun W. Navigating phase diagram complexity to guide robotic inorganic materials synthesis. *Nat Synth.* 2024;3(5):606-14.

<https://doi.org/10.1038/s44160-024-00502-y>.

Bayley O, Savino E, Slattery A, Noël T. Autonomous chemistry: Navigating self-driving labs in chemical and material sciences. *Matter*. 2024;7(7):2382-98.

<https://doi.org/10.1016/j.matt.2024.06.003>.

Dave A, Mitchell J, Burke S, Lin H, Whitacre J, Viswanathan V. Autonomous optimization of non-aqueous Li-ion battery electrolytes via robotic experimentation and machine learning coupling. *Nat Commun*. 2022;13(1):5454.

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

Bai J, Mosbach S, Taylor CJ, Karan D, Lee KF, Rihm SD, et al. A dynamic knowledge graph approach to distributed self-driving laboratories. *Nat Commun*. 2024;15(1):462.

<https://doi.org/10.1038/s41467-023-44599-9>.

Statt MJ, Rohr BA, Guevarra D, Breeden J, Suram SK, Gregoire JM. The materials experiment knowledge graph. *Digit Discov*. 2023;2(4):909-14.

<https://doi.org/10.1039/D3DD00067B>.

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

<https://doi.org/10.1016/j.matt.2024.01.005>.