Peng *et al. AI Agent.* **2025**, *1*, 5
DOI: 10.20517/aiagent.2025.07

AI Agent

Commentary

Open Access

Check for updates

Accelerating multimetallic catalyst discovery with robotics and agentic AI

Jiayu Peng[1], Chuanyu Liu[1], Yiwen Luo[1], Kritarth Dandapat[2]

[1]Department of Materials Design and Innovation, University at Buffalo, Buffalo, NY 14260, USA.
[2]Department of Computer Science and Engineering, University at Buffalo, Buffalo, NY 14260, USA.

Correspondence to: Prof. Jiayu Peng, Department of Materials Design and Innovation, University at Buffalo, Buffalo, NY 14260, USA. E-mail: jypeng@buffalo.edu

How to cite this article: Peng, J.; Liu, C.; Luo, Y.; Dandapat, K. Accelerating multimetallic catalyst discovery with robotics and agentic AI. *AI Agent* **2025**, *1*, 5. https://dx.doi.org/10.20517/aiagent.2025.07

Received: 20 Oct 2025 **First Decision:** 30 Oct 2025 **Revised:** 2 Nov 2025 **Accepted:** 4 Nov 2025 **Published:** 24 Nov 2025

Academic Editor: Aloysius Soon **Copy Editor:** Xing-Yue Zhang **Production Editor:** Xing-Yue Zhang

Abstract

The design space of catalyst materials spans composition, processing, atomistic structure, and microstructure. As materials become more complex, the dimensionality of this parameter space for catalyst design grows combinatorially. Conventional active learning approaches operate on a single data stream and stay decoupled from the messy reality of experiments, limiting their efficiency and reproducibility in real-world catalyst optimization. To tackle this limitation, in a recent issue of *Nature*, Li *et al.* developed a robotic platform, Copilot for Real-world Experimental Scientists (CRESt), which facilitates multimetallic catalyst discovery in a multiplex parameter space by combining multimodal large vision-language models, knowledge-assisted Bayesian optimization, and robotic automation of synthesis, characterization, and electrochemical tests. Deployed on a direct formate fuel cell use case, CRESt efficiently explored hundreds of compositions and thousands of tests in months to deliver an octonary multimetallic electrocatalyst with excellent device-level performance at reduced noble-metal loading. In this Commentary, we highlight CRESt's technical merits, while also outlining a forward agenda to translate systems such as CRESt from proof-of-concept, bespoke demonstrations to widely adoptable, scientifically robust agentic artificial intelligence for self-driving laboratories.

Keywords: AI for science, catalysis, electrochemistry, self-driving labs, vision language model, Bayesian optimization, multimodal machine learning

Electrocatalysis sits at the core of decarbonizing various industrial sectors by electrifying the production and utilization of green chemicals and clean fuels through electrochemical energy conversion and storage technologies, e.g., electrolysis and fuel cells[1]. Elucidating the reactivity[2] and transformation[3] of material surfaces while balancing catalytic activity[4] and stability[5] remains a scientific challenge to fully unlock the promise of electrocatalysis for various sustainable energy technologies. Notably, the development of electrocatalysts has been especially slow for the electrochemical oxidation of liquid fuels (e.g., methanol, ethanol, and formate)[6]. As an example, the electro-oxidation of formate is highly difficult, which accounts for over half of the energy losses in the direct formate fuel cell (DFFC) owing to its complex reaction pathways, hard-to-tune surface intermediate interactions, and parasitic by-products[6].

Multimetallic catalysts, e.g., high-entropy alloys (HEAs), have recently been shown to exhibit potential for complex reactions such as formate oxidation by providing chemically and structurally diverse surface sites that lead to more optimized reaction energetics than simpler catalysts, e.g., pure metals and binary alloys[7]. In an ideal HEA solid solution, all its constituent elements are randomly distributed to maximize the configurational entropy. This randomness yields a rich ensemble of surface metal sites and, consequently, a quasi-continuous distribution of adsorption energetics for reaction intermediates. Such binding environments can drive precise optimization of local electronic structure and site-specific reaction energetics via strain and ligand effects, thereby enabling the possibility to design surface motifs and discover novel materials with near-optimal catalytic performance.

Unfortunately, multimetallic catalyst optimization has been notoriously difficult due to the vastness of the underlying parameter space and the elusive, convoluted correlations between various design variables. Intrinsic parameters, including chemical composition, atomic structure, and microstructure, are often correlated non-linearly[8], which are further confounded by extrinsic factors, such as synthetic, processing, and operating conditions[9]. Active learning driven by Bayesian optimization (BO) with the Gaussian process has emerged as a principled way to balance exploration and exploitation in a complex parameter space[10], boosting catalyst discovery compared with traditional ad-hoc or exhaustive search[11]. BO has recently been demonstrated to traverse multimetallic spaces efficiently by steering experiments toward promising regions while quantifying uncertainty[12] - an approach that complements high-throughput materials synthesis, characterization, and performance testing[13-17]. However, such BO implementations have been unimodal, as they optimize over a single data stream, e.g., chemical composition or processing parameters. Thus, a missed opportunity resides in multimodal active learning, where an agentic artificial intelligence (AI) system can be used to ingest literature knowledge, text, images, and experiment logs, such that the search policy is informed by what human scientists actually use, with the realistic complexity in materials design[18].

To bridge this technical gap between AI agents and self-driving laboratories, Li *et al.* established the Copilot for Real-world Experimental Scientists (CRESt) AI system[19], which integrates multimodal large vision-language models (LVLMs) and knowledge-assisted active learning with robotics that perform high-throughput synthesis, characterization, and electrochemistry at scale [Figure 1]. CRESt is a three-part platform: (1) a natural-language user interface; (2) an LVLM-driven BO back-end that embeds text, composition, and images into a reduced BO search space; and (3) actuators for automated catalyst preparation, characterization, and testing. In practice, researchers can "chat" with the system, which synthesizes materials, performs measurements, monitors experimental runs with cameras and LVLMs, suggests the next candidates to assess, and proposes fixes when anomalies arise - better mirroring how expert materials researchers actually reason, in contrast to conventional BO algorithms that operate on a unimodal data stream.

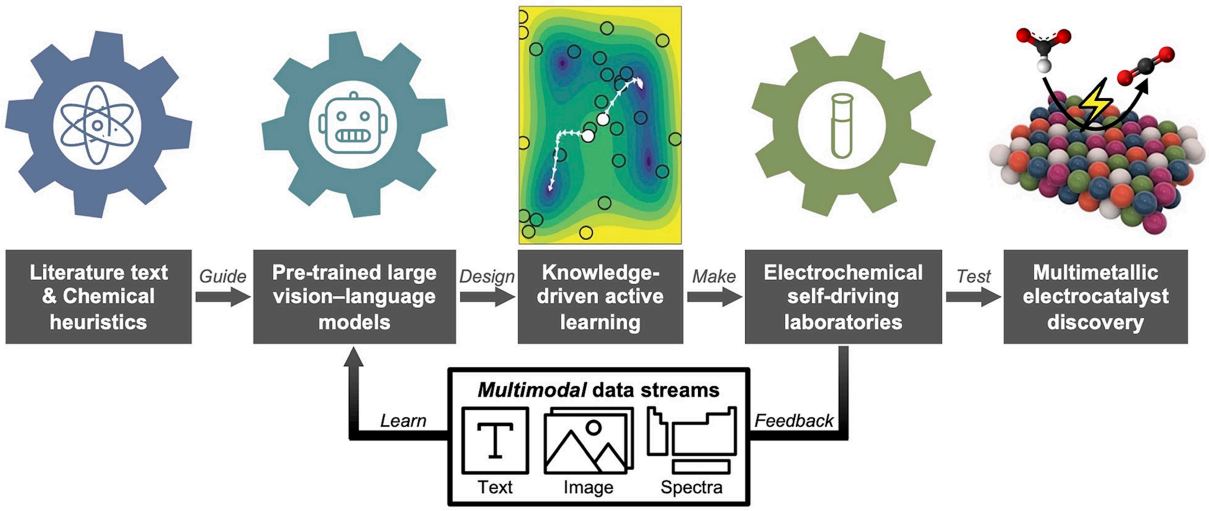

Figure 1. Schematic illustration of CRESt, which utilizes an LVLM-powered agentic AI system to embed diverse information (such as literature text, microstructural images, and experimental spectra) into a reduced search space, couple it to active learning, and further streamlines high-throughput synthesis, characterization, and tests for catalyst discovery. CRESt: Copilot for Real-world Experimental Scientists; LVLM: large vision-language model.

An algorithmic innovation of CRESt is its LVLM-powered active learning approach that effectively incorporates knowledge-embedding-based search space reduction into an adaptive exploration-exploitation strategy. CRESt embeds prior knowledge by querying and encoding prior literature related to the reaction of interest into text embeddings, while also extracting image-derived morphological descriptors using authentic characterization data from autonomous scanning electron microscopy. These text and image embeddings are compressed into a small set of latent variables that preserve most of their variance via principal component analysis. BO with a knowledge-gradient acquisition function[20] then proposes the next chemical compositions to investigate, mapping measured performance back to formulation. CRESt can adapt the exploration-exploitation trade-off dynamically without manual tuning of this balance, achieved by introducing a Lagrangian multiplier[21] motivated by policy-improvement constraints often found in reinforcement learning.

For the hardware, CRESt coordinates liquid handling, carbothermal shock synthesis, microscopy, diffraction, and electrochemistry, orchestrating all experimental subsystems to close the end-to-end optimization loop at the precise cadence that agentic AI demands. High-throughput carbothermal shock synthesis was used due to its simplicity to automate and its benchmarked efficacy to produce high-quality single-phase HEA nanoparticles[22]. Applied to formate oxidation with the oxygen reduction reaction at the counter electrode in a DFFC, CRESt synthesized over 900 chemistries and ran ~3,500 electrochemical tests in three months, a pace that is challenging to achieve manually[19]. This campaign surfaced an octonary HEA catalyst within the Pd-Pt-Cu-Au-Ir-Ce-Nb-Cr multimetallic space with a 9.3-fold improvement in cost-specific performance relative to conventional Pd catalysts. Such an optimized recipe delivered high experimental power density in a working DFFC at only one-quarter the precious-metal loading typical of these energy conversion devices. Interestingly, CRESt's monitoring and self-diagnosis can matter as much as optimization. Camera streams with LVLMs were found to detect subtle deviations and suggest possible corrections, improving reproducibility while keeping human researchers in the loop.

This proof-of-principle study demonstrates the promise of combining multimodal AI agents with self-driving laboratories to boost the discovery and optimization of materials or molecules across a massive design space and account for the real-world complexity of experimental variables and data streams. Similar

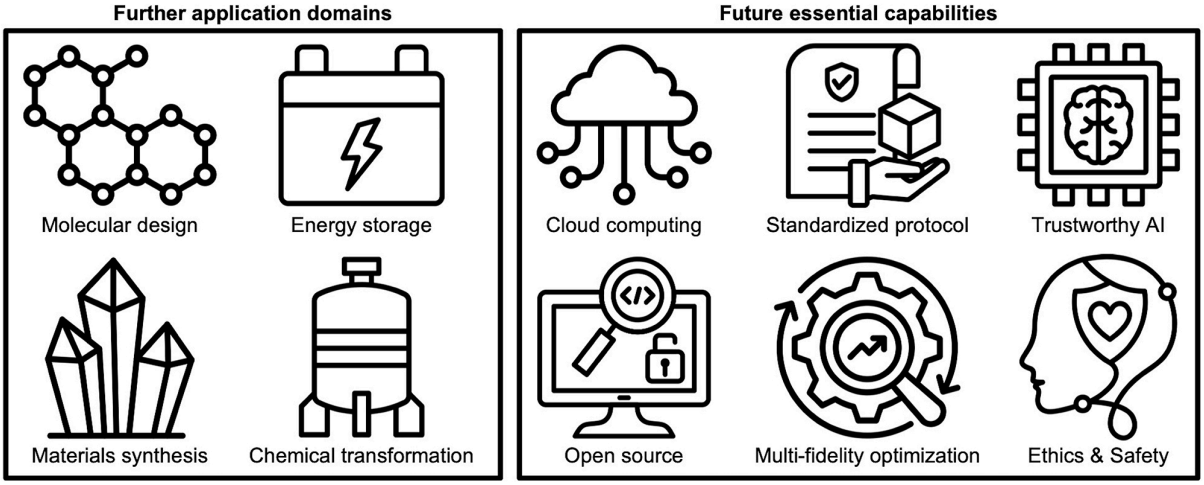

Figure 2. Possible extensions to other research directions and further advances that need to be accomplished to unleash the full potential of agentic AI for self-driving laboratories. AI: Artificial intelligence.

agentic AI frameworks driven by large language models (LLMs) have been demonstrated for other self-driving laboratories, but these representative examples have been mostly limited to molecular synthesis[23-26]. It is intriguing to explore if AI agents can be applied to other research directions [Figure 2]. In thermocatalysis, a future route is to integrate robotic reactors with inline analytics and use microscopy or spectroscopy as feedback in an active learning loop. Such a setup may autonomously navigate the space of HEA nanoparticles for selective hydrogenation[27] or oxidative dehydrogenation[28], where the interplay between composition, microstructure, and surface states challenges traditional manual search. Relatedly, large, public databases, e.g., Digital Catalysis Platform (DigCat: https://www.digcat.org/)[29], offer cloud-based AI tools for catalyst design with closed-loop workflows, complementing lab-in-the-loop systems, including CRESt. For energy storage, this framework naturally extends to electrolyte formulation[30] and electrode synthesis optimization[31], e.g., by fusing literature text, microscopy, and time-series electrochemical data to triage long-horizon cycling experiments. To boost materials synthesis, coupling state-of-the-art flow chemistry platforms[32] with LVLMs can co-optimize precursors, reagents, and conditions and use text-mined precedents and online spectroscopies to propose fixes when runs drift. Across various application domains, this central concept - compressing multimodal prior knowledge into a latent design space for actively learning with robotics - offers a general path to transform heterogeneous evidence into faster, more reliable materials discovery.

It is also essential to emphasize that several limitations still prevail in today's agentic AI frameworks for materials and molecular design [Figure 2]. To begin with, they remain costly bespoke systems with significant engineering overhead. Future progress will hinge on modular, open-source hardware design and cloud-based orchestration stacks that may lower the integration burden and democratize the use of these self-driving laboratories[33]. Similarly, the generalizability of such autonomous systems needs to be boosted. To move from single-use exemplars to reconfigurable laboratories, these high-throughput systems should allow plug-and-play task flexibilities and protocol abstractions, such that the same copilot can pivot from one research topic to another without a comprehensive redesign[34]. Concretely, emerging open standards for tool-data connectivity (e.g., the Model Context Protocol)[35] and tools for autonomous experiments (e.g., ChemOS[36,37] and AlabOS[38], two recently developed orchestrating architecture and workflow management frameworks) are likely to boost further improvements in the reconfigurability of self-driving laboratories. It is equally essential to note that scaling agentic AI alone does not resolve all limitations of machine learning. Models trained on historical datasets can degrade under distribution shift and struggle to extrapolate to truly novel chemistries or structures, necessitating out-of-distribution (OOD)-aware evaluation and model

selection[39]. Advances in underlying methods, e.g., uncertainty quantification[40] and OOD detection[41], as well as multi-fidelity BO[42], remain vital in facilitating reliable self-driving laboratories for materials discovery. Moreover, several challenges have recently been observed in LLM- and LVLM-driven AI agents, e.g., data biases, knowledge limitations, hallucinations, and poor model interpretability[43,44]. Changes in current evaluation schemes, such as rewarding calibrated abstention rather than guessing[45], might suppress hallucinations by informing abstention-aware planning and promoting retrieval-augmented decision policies in agentic AI loops. Additionally, high-quality, multimodal datasets with well-defined tasks and benchmarks are still urgently required in order to systematically evaluate the chemical knowledge and reasoning abilities of LLMs and LVLMs[46] and improve the robustness and transparency of the resulting AI agents in driving real-world materials discovery[47]. Finally, ethical and safety considerations remain crucial for both self-driving laboratories[48] and agentic AI[49]. For instance, ensuring that hazardous experiments are avoided is paramount, which could be achieved by explicitly leveraging watchdog AIs that vet proposed experiments against hazard databases and chemical safety rules. Likewise, the introduction of multimodal AI agents into self-driving laboratories presents ethical dimensions that need to be managed, including preventing the misuse of such technologies toward harmful ends, e.g., creating toxic chemicals or designing chemical weapon formulations[50]. With these considerations and future advances - more robust standardization, better data foundation, more rigorous evaluation, safer autonomy - multimodal, agentic AI will continue to evolve from proof-of-concept demonstrations such as CRESt into widely adoptable scientific copilots for every researcher to close the design-make-test-learn loop for materials innovation at scale.

DECLARATIONS

Authors' contributions

Conceptualization (lead), funding acquisition (lead), investigation (lead), writing - original draft (lead), and writing - review & editing (lead): Peng, J.

Writing - review & editing (supporting): Liu, C.; Luo, Y.; Dandapat, K.

Availability of data and materials

Not applicable.

Financial support and sponsorship

Peng, J. acknowledges the University at Buffalo for start-up funds to support this work.

Conflicts of interest

All authors declared that there are no conflicts of interest.

Ethical approval and consent to participate

Not applicable.

Consent for publication

Not applicable.

REFERENCES

1. Wei, C.; Rao, R. R.; Peng, J.; et al. Recommended practices and benchmark activity for hydrogen and oxygen electrocatalysis in water splitting and fuel cells. *Adv. Mater.* **2019**, *31*, 1806296. DOI PubMed
2. Wang, T.; Iriawan, H.; Peng, J.; et al. Confined water for catalysis: thermodynamic properties and reaction kinetics. *Chem. Rev.* **2025**, *125*, 1420-67. DOI PubMed

3. Peng, J. Toward data-driven predictive modeling of electrocatalyst stability and surface reconstruction. *J. Chem. Phys.* **2025**, *163*, 040902. DOI
4. Yuan, S.; Peng, J.; Cai, B.; et al. Tunable metal hydroxide-organic frameworks for catalysing oxygen evolution. *Nat. Mater.* **2022**, *21*, 673-80. DOI PubMed
5. Peng, J.; Giner-sanz, J. J.; Giordano, L.; et al. Design principles for transition metal nitride stability and ammonia generation in acid. *Joule.* **2023**, *7*, 150-67. DOI
6. Zheng, D. J.; Peng, J.; Mccormack, K.; et al. Uniting activity design principles of anode catalysts for direct liquid fuel cells. *EES. Catal.* **2024**, *2*, 1186-209. DOI
7. Löffler, T.; Ludwig, A.; Rossmeisl, J.; Schuhmann, W. What makes high-entropy alloys exceptional electrocatalysts? *Angew. Chem. Int. Ed. Engl.* **2021**, *60*, 26894-903. DOI PubMed PMC
8. Hsu, W. L.; Tsai, C. W.; Yeh, A. C.; Yeh, J. W. Clarifying the four core effects of high-entropy materials. *Nat. Rev. Chem.* **2024**, *8*, 471-85. DOI PubMed
9. Kar, N.; Skrabalak, S. E. Synthetic methods for high-entropy nanomaterials. *Nat. Rev. Mater.* **2025**, *10*, 638-53. DOI
10. Peng, J.; Schwalbe-koda, D.; Akkiraju, K.; et al. Human- and machine-centred designs of molecules and materials for sustainability and decarbonization. *Nat. Rev. Mater.* **2022**, *7*, 991-1009. DOI
11. Peng, J.; Damewood, J. K.; Karaguesian, J.; Gómez-bombarelli, R.; Shao-horn, Y. Navigating multimetallic catalyst space with Bayesian optimization. *Joule.* **2021**, *5*, 3069-71. DOI
12. Rohr, B.; Stein, H. S.; Guevarra, D.; et al. Benchmarking the acceleration of materials discovery by sequential learning. *Chem. Sci.* **2020**, *11*, 2696-706. DOI PubMed PMC
13. Mints, V. A.; Pedersen, J. K.; Bagger, A.; et al. Exploring the composition space of high-entropy alloy nanoparticles for the electrocatalytic H_2/CO oxidation with Bayesian optimization. *ACS. Catal.* **2022**, *12*, 11263-71. DOI
14. Jenewein, K. J.; Torresi, L.; Haghmoradi, N.; Kormányos, A.; Friederich, P.; Cherevko, S. Navigating the unknown with AI: multiobjective Bayesian optimization of non-noble acidic OER catalysts. *J. Mater. Chem. A.* **2024**, *12*, 3072-83. DOI
15. Esterhuizen, J. A.; Mathur, A.; Goldsmith, B. R.; Linic, S. High-performance iridium-molybdenum oxide electrocatalysts for water oxidation in acid: Bayesian optimization discovery and experimental testing. *J. Am. Chem. Soc.* **2024**, *146*, 5511-22. DOI PubMed
16. Tadgell, C. A.; Kato, M.; Dieb, S.; et al. Machine learning-assisted development of platinum-free RuNiCo nanocages for electrocatalytic hydrogen oxidation reaction in acidic media. *ACS. Appl. Energy. Mater.* **2025**, *8*, 14052-7. DOI
17. Niu, X.; Chen, Y.; Sun, M.; et al. Bayesian learning-assisted catalyst discovery for efficient iridium utilization in electrochemical water splitting. *Sci. Adv.* **2025**, *11*, eadw0894. DOI PubMed PMC
18. Xin, H.; Kitchin, J.; López, N.; et al. Roadmap for transforming heterogeneous catalysis with artificial intelligence. *ChemRxiv* **2025**. Available online: https://doi.org/10.26434/chemrxiv-2025-chn4j (accessed 20 November 2025).
19. Zhang, Z.; Ren, Z.; Hsu, C. W.; et al. A multimodal robotic platform for multi-element electrocatalyst discovery. *Nature.* **2025**, *647*, 390-6. DOI PubMed
20. Frazier, P. I.; Powell, W. B.; Dayanik, S. A knowledge-gradient policy for sequential information collection. *SIAM. J. Control. Optim.* **2008**, *47*, 2410-39. DOI
21. Chen, E. Hong, Z.W.; Pajarinen, J.; Agrawal, P. Redeeming intrinsic rewards via constrained optimization. In *Advances in Neural Information Processing Systems 35 (NeurIPS 2022)*, New Orleans, LA, USA, and online, November 28-December 9, 2022; Koyejo, S., Mohamed, S., Agarwal, A., Belgrave, D., Cho, K., Oh, A., Eds.; Curran Associates, Inc.: Red Hook, NY, USA, 2022; pp 4996-5008. DOI
22. Yao, Y.; Huang, Z.; Xie, P.; et al. Carbothermal shock synthesis of high-entropy-alloy nanoparticles. *Science.* **2018**, *359*, 1489-94. DOI PubMed
23. Boiko, D. A.; Macknight, R.; Kline, B.; Gomes, G. Autonomous chemical research with large language models. *Nature.* **2023**, *624*, 570-8. DOI PubMed PMC
24. M. Bran, A.; Cox, S.; Schilter, O.; Baldassari, C.; White, A. D.; Schwaller, P. Augmenting large language models with chemistry tools. *Nat. Mach. Intell.* **2024**, *6*, 525-35. DOI PubMed PMC
25. Darvish, K.; Skreta, M.; Zhao, Y.; et al. ORGANA: a robotic assistant for automated chemistry experimentation and characterization. *Matter.* **2025**, *8*, 101897. DOI
26. Song, T.; Luo, M.; Zhang, X.; et al. A multiagent-driven robotic AI chemist enabling autonomous chemical research on demand. *J. Am. Chem. Soc.* **2025**, *147*, 12534-45. DOI PubMed
27. Jing, Z.; Guo, Y.; Wang, Q.; et al. Ambient hydrogenation of solid aromatics enabled by a high entropy alloy nanocatalyst. *Nat. Commun.* **2024**, *15*, 5806. DOI PubMed PMC
28. Luo, J.; Li, X.; Ye, Y.; et al. Progressive fabrication of a Pt-based high-entropy-alloy catalyst toward highly efficient propane dehydrogenation. *Angew. Chem. Int. Ed. Engl.* **2024**, *64*, e202419093. DOI PubMed
29. Zhang, D.; Li, H. Digital Catalysis Platform (DigCat): a gateway to big data and AI-powered innovations in catalysis. *ChemRxiv* **2024**. Available online: https://doi.org/10.26434/chemrxiv-2024-9lpb9 (accessed 20 November 2025).
30. Stolberg, M. A.; Lopez, J.; Cawthern, S. D.; et al. A data-driven platform for automated characterization of polymer electrolytes. *Matter.* **2025**, *8*, 102129. DOI
31. Szymanski, N. J.; Rendy, B.; Fei, Y.; et al. An autonomous laboratory for the accelerated synthesis of novel materials. *Nature.* **2023**, *624*, 86-91. DOI PubMed PMC
32. Delgado-licona, F.; Alsaiari, A.; Dickerson, H.; et al. Flow-driven data intensification to accelerate autonomous inorganic materials discovery. *Nat. Chem. Eng.* **2025**, *2*, 436-46. DOI
33. Abolhasani, M.; Kumacheva, E. The rise of self-driving labs in chemical and materials sciences. *Nat. Synth.* **2023**, *2*, 483-92. DOI
34. Tom, G.; Schmid, S. P.; Baird, S. G.; et al. Self-driving laboratories for chemistry and materials science. *Chem. Rev.* **2024**, *124*, 9633-732. DOI PubMed PMC
35. Anthropic. Introducing the model context protocol. November 24, 2024. https://www.anthropic.com/news/model-context-protocol (accessed 2025-11-20).
36. Roch, L. M.; Häse, F.; Kreisbeck, C.; et al. ChemOS: an orchestration software to democratize autonomous discovery. *PLoS. ONE.* **2020**, *15*, e0229862. DOI PubMed PMC
37. Sim, M.; Vakili, M. G.; Strieth-kalthoff, F.; et al. ChemOS2.0: an orchestration architecture for chemical self-driving laboratories. *Matter.* **2024**, *7*, 2959-77. DOI
38. Fei, Y.; Rendy, B.; Kumar, R.; et al. AlabOS: a Python-based reconfigurable workflow management framework for autonomous laboratories. *Digit. Discov.* **2024**, *3*, 2275-88. DOI

39. Li, K.; Decost, B.; Choudhary, K.; Greenwood, M.; Hattrick-simpers, J. A critical examination of robustness and generalizability of machine learning prediction of materials properties. *npj. Comput. Mater.* **2023**, *9*, 55. DOI
40. Dai, J.; Adhikari, S.; Wen, M. Uncertainty quantification and propagation in atomistic machine learning. *Rev. Chem. Eng.* **2025**, *41*, 333-57. DOI
41. Li, K.; Rubungo, A. N.; Lei, X.; et al. Probing out-of-distribution generalization in machine learning for materials. *Commun. Mater.* **2025**, *6*, 9. DOI
42. Sabanza-gil, V.; Barbano, R.; Pacheco Gutiérrez, D.; et al. Best practices for multi-fidelity Bayesian optimization in materials and molecular research. *Nat. Comput. Sci.* **2025**, *5*, 572-81. DOI PubMed
43. Balcells, D. Co-intelligent design of catalysis research with large language models: hype or reality? *ACS. Catal.* **2025**, *15*, 16412-20. DOI
44. Xin, H.; Kitchin, J. R.; Kulik, H. J. Towards agentic science for advancing scientific discovery. *Nat. Mach. Intell.* **2025**, *7*, 1373-5. DOI
45. Kalai, A. T.; Nachum, O.; Vempala, S. S.; Zhang, E. Why language models hallucinate. *arXiv* **2025**, arXiv:2509.04664. Available online: https://doi.org/10.48550/arXiv.2509.04664 (accessed 20 November 2025).
46. Mirza, A.; Alampara, N.; Kunchapu, S.; et al. A framework for evaluating the chemical knowledge and reasoning abilities of large language models against the expertise of chemists. *Nature. Chem.* **2025**, *17*, 1027-34. DOI PubMed PMC
47. Miret, S.; Krishnan, N. M. A. Enabling large language models for real-world materials discovery. *Nat. Mach. Intell.* **2025**, *7*, 991-8. DOI
48. Leong, S. X.; Griesbach, C. E.; Zhang, R.; et al. Steering towards safe self-driving laboratories. *Nat. Rev. Chem.* **2025**, *9*, 707-22. DOI PubMed
49. Macknight, R.; Boiko, D. A.; Regio, J. E.; Gallegos, L. C.; Neukomm, T. A.; Gomes, G. Rethinking chemical research in the age of large language models. *Nat. Comput. Sci.* **2025**, *5*, 715-26. DOI PubMed
50. Spotte-smith, E. W. C. Considering the ethics of large machine learning models in the chemical sciences. *Mach. Learn.:. Sci. Technol.* **2025**, *6*, 035007. DOI