Materials Horizons

ROYAL SOCIETY OF CHEMISTRY

COMMUNICATION

View Article Online
View Journal



Cite this: DOI: 10.1039/d6mh00475j

Received 12th March 2026,
Accepted 5th May 2026

DOI: 10.1039/d6mh00475j

rsc.li/materials-horizons

# Agentic workflow enables the recovery of critical materials from complex feedstocks *via* selective precipitation

Andrew Ritchhar, †[a] Sarah I. Allec, †[a] Pravalika Butreddy, [a] Krista Kulesa, [a] Qingpu Wang, [b] Dan Thien Nguyen, [a] Maxim Ziatdinov *[a] and Elias Nakouzi *[ac]



**We present a multi-agentic workflow for critical materials recovery that deploys a series of AI agents and automated instruments to recover critical materials from produced water and magnet leachates. This approach achieves selective precipitation from real-world feedstocks using simple chemicals, accelerating the optimization of efficient, adaptable, and scalable separations to a timeline of days, rather than months and years.**

The recovery of critical materials from complex feedstocks is a problem of increasing urgency due to the surging demand for these elements and their importance to energy technologies.[1] One of the cornerstone techniques for critical materials separation is selective precipitation, wherein a chemical reagent added to the mixture reacts with the less soluble ions, creating solid precipitates that can subsequently be recovered. Precipitation is widely applied in the separations industry due to its scalability, chemical versatility, and low energy expense.[2] Recent studies have extended this approach to unconventional feedstocks, including polymetallic nodules, recycled electronics, seawater, low-grade ores, and more.[3–10] The state-of-the-art research in this field focuses on developing tailored and engineered materials such as membranes, metal organic frameworks, ionic liquids, and ligands, with highly optimized properties such as pore size, binding energy, or ion exchange coefficients.[11,12] These techniques are particularly viable when target ions are dilute, high purity is required, and competing ions have similar chemistry. However, selective precipitation offers a straightforward option with simple commodity chemicals that perform reasonably well, especially when factoring energy-efficiency and economic considerations.[13]

One key limitation of selective precipitation is that the products often incorporate significant amounts of impurities, resulting in inclusions or solid solutions that require further purification. This challenge is compounded by the nature of real-world feedstocks, which are typically highly variable and complex, with compositions that are highly source and method dependant, often including multiple competing ions with similar physicochemical properties.[14] Predicting the composition of the precipitates from complex feedstocks and thus the separation efficiency remains beyond the limits of physics-based speciation and crystal growth models. As such, optimizing precipitation-based separations across this vast parameter space presents an excellent use case for artificial intelligence (AI)-powered workflows and autonomous experimentation.

## New concepts

The recovery of critical materials from complex feedstocks is a problem of increasing urgency in the materials research community. Considerable efforts aim at expanding and optimizing hydrometallurgy methods, due to the complexity and variability of real-world feedstocks. In this work, we present a multi-agentic workflow, CICERO (Computer Intelligence for Critical Element Recovery and Optimization), powered by the in-house-developed SciLink platform, that leverages automated chemical synthesis and characterization instruments, to optimize separation recipes within days rather than months or years. We tested CICERO with three types of feedstocks and achieved successful critical material separations of magnesium from produced water, samarium from leached SmCo magnets, and neodymium/praseodymium from leached NdFeB magnets. CICERO uses as input an initial characterization of the feedstock composition, augmented with supporting documents, to evaluate the feedstock, hypothesize a separations pathway, design an experimental campaign, generate scripts for instrument operation, execute the experiments, process the results, and propose additional experiments to converge towards optimal % yield and % purity of the target critical materials. This work presents a new approach for optimizing critical materials separations and demonstrates how locally-developed, domain-specific agentic workflows can enhance productivity and efficiency in solving materials science and engineering problems.

[a] *Physical and Computational Sciences Directorate, Pacific Northwest National Laboratory, Richland, Washington 99354, USA. E-mail: maxim.ziatdinov@pnnl.gov, elias.nakouzi@pnnl.gov*
[b] *Physical and Computational Sciences Directorate, Pacific Northwest National Laboratory, Seattle, Washington 98109, USA*
[c] *Department of Chemical Engineering, University of Washington, Seattle, Washington 98195, USA*
† These authors are co-first authors who have contributed equally to this work.

Recent advances in artificial intelligence,[15–19] particularly large language models (LLMs),[20–22] have enhanced and widened the possibilities for AI-driven scientific experimentation. However, effectively deploying these capabilities in laboratory settings remains challenging. Stand-alone LLMs lack the specialization and niche domain knowledge required for many areas of scientific research. Conversely, platforms built specifically for laboratory automation – including LLM-enabled systems – often suffer from being overly rigid, demanding extensive integration efforts and programming expertise to customize and deploy. Truly transformative AI-driven scientific innovation will require platforms that are modular and adaptive yet resource-efficient, while remaining accessible to domain scientists. Agentic AI – systems in which LLM-based agents can autonomously plan, execute, and iteratively refine multi-step workflows – offers a promising path forward by combining the flexible reasoning of LLMs with the ability to dynamically invoke specialized tools, instruments, and domain knowledge as needed, all primarily through natural language interaction.

In realizing this vision, one of the biggest challenges is defining scientific problems that are simultaneously ambitious enough to merit the development of agentic workflows, while being tractable with current AI agents, lab instrumentation, and reasonable constraints on the science questions being asked. In this context, adapting selective precipitation using simple commodity chemicals and existing liquid handling automation is an impactful and attainable goal. Here, we demonstrate an agentic AI-guided closed-loop workflow that autonomously plans, evaluates, and optimizes selective precipitation campaigns for critical materials recovery from complex feedstocks and waste streams. The goal of this experimental campaign is not to develop new knowledge or design state-of-the-art materials for chemical separations, but rather to optimize separations from complex real-world feedstocks using simple and scalable precipitation reactions.

Our workflow is built on SciLink,[23] a recently developed LLM-powered multi-agent platform that provides a modular library of plug-and-play agents for automating experimental design, data analysis, and computational modelling tasks in materials research. In particular, we developed and integrated into SciLink a planning subsystem for critical materials recovery, which guides researchers through the full campaign loop from literature-grounded hypothesis generation and economic viability screening, through automated data analysis, to closed-loop Bayesian optimization; a workflow we refer to as CICERO (computer intelligence for critical elements recovery and optimization). Within CICERO, the planning subsystem manages the full lifecycle of an experimental campaign through three coordinated SciLink sub-agents. The Planning Agent ingests domain literature (papers, technical reports) and any existing experimental data into a retrieval-augmented generation pipeline, producing testable hypotheses, detailed experimental protocols, and economic viability screenings – all grounded in and traceable to the source documents. Critically, the viability screening step acts as an economic filter that narrows the hypothesis and planning space before any experiments begin, steering the campaign toward materials, processes, and parameter ranges worth pursuing given real-world cost drivers, market viability, and scalability constraints. Once experiments are underway, raw instrument data is converted into scalar descriptors suitable for mathematical optimization. While SciLink includes a Scalarizer Agent to automate this task by auto-generating analysis scripts, the data conversion in this work was straightforward, and we employed a pre-existing, validated script. These descriptors then feed into the Bayesian optimization (BO) Agent, which drives a Bayesian optimization loop[24] – supporting single- and multi-objective campaigns, budget-aware acquisition strategies, batch parallelism for high-throughput platforms, and physical constraint handling for real-world setups such as multi-well plates.

We tested CICERO by processing three types of feedstocks: (1) produced water from oil and gas extraction in Oklahoma, (2) commercially sourced samarium–cobalt (SmCo) permanent magnets, and (3) commercially sourced neodymium–iron–boron (NdFeB) permanent magnets. Details on the feedstock compositions and pre-processing steps, such as magnet acid leaching, are provided in the SI. For each feedstock, the end point after processing the sample with CICERO is a well-defined, feedstock-specific protocol for the recovery of critical materials using simple commodity chemicals. The goal of this work is not to develop new separation chemistries and technologies, but rather to rapidly identify and optimize separation methodologies that are feedstock-specific, economically feasible, and readily scalable as they rely exclusively on commodity chemicals already used in industrial separations.

In the following, we provide additional details on the sequential steps, AI agents, and current roles of humans in CICERO. (1) The initial "Diagnose" step includes characterization of the feedstock chemistry, providing the basis data for attempting hundreds of separation experiments in the subsequent steps. Specifically, quantitative Inductively Coupled Plasma Mass Spectrometry (ICP-MS) measurements produce a table of the identity and concentrations of the elements present in the feedstock. The SciLink Planning Agent ingests this literature and data to map the problem space. (2) Next is the "Evaluate" step, where the SciLink Planning Agent determines the target critical materials in the feedstock based on a preliminary technoeconomic analysis of value, concentration, criticality, and anticipated range of product purity. The agent is optionally augmented at run time with additional data, for example in our described use case including US DOE Critical Materials Assessment documents. (3) The decision on the target critical materials is passed onto the "Hypothesize" step, where the SciLink Planning Agent generates one or more testable pathways for recovering the target critical material. By prompting, we specifically target in this work a series of precipitation and dissolution steps; an approach that remains the backbone of the chemical separations industry. To this end, the agent is constrained to using commodity chemicals, such as simple acids and bases, as well as precipitants, such as oxalic acid, citric acid, sodium bicarbonate, or hydrogen peroxide.

The next steps involve physical experiments: (4) In the "Experiment" step, the agent converts the hypothesis into an initial batch of 96 experiments that span the relevant parameter space, including the identity and concentration of the chemicals

 

used for separation, as well as the sequence and time of adding these chemicals. The agent then provides an actionable Python script that automatically executes the experiments in a 96-well plate using an Opentrons robot. After the experiments are complete, a human operator moves the samples between the Opentrons robot and a microplate centrifuge for multiple cycles of washing and decanting, before using automatically generated code to dissolve the purified products in preparation for ICP-MS measurements. A human then moves the plate to the ICP-MS instrument, which automatically measures the chemical composition in the products of all the samples. (5) The results are then passed to a pre-existing script for processing the ICP-MS data files and resolving the composition of the products across the parameter space. (6) The final "Refine" step involves determining whether satisfactory separations have been achieved by evaluating the % purity by mol relative to the other metals and % yield of the target critical material. This involves the optional use of the Bayesian Optimization agent or re-prompting the planning agent to recommend additional iterations or other targets for new experiments, closing the loop and converging towards an optimized separations method. CICERO is programmed to offer a choice of either fully automated analysis and planning or human-in-the-loop feedback.

In the experiments described, human feedback was kept minimal and used at three points: in the experimental agent to modify maximum allowable reagent concentrations, in code generation to supervise and fix any minor bugs, and in refinement to inform the target purity metric of subsequent steps. All agents utilized Google's gemini-2.5-pro model.[25] Verbatim SciLink agent input and output text are available in the SI.

Using CICERO, we processed three types of real-world feedstocks, deploying distinct approaches for critical materials separations as illustrated in the multi-panel Fig. 2. For the produced water feedstock, CICERO identified magnesium as the highest value target. Despite noting the presence of other possibly more valuable components, such as copper on a per-atom basis, magnesium was selected due to its exceedingly high concentration (1840 ppm) and cost effectiveness for recovery. The agent-generated hypothesis was tuning the pH to alkaline conditions to selectively precipitate magnesium hydroxide, bypassing the other alkaline earth metals present in substantial amounts, including calcium (>13 000 ppm), strontium (450 ppm), and barium (2 ppm), as well as more soluble ions such as sodium (>6000 ppm). This approach involved the addition of a single precipitating agent, sodium hydroxide, in various concentrations and evaluating the composition of the solid precipitates. The agent opted for a molar ratio of NaOH to Mg between 1.5 : 1 and 2 : 1, targeting 8 experimental conditions with 12 sample duplicates. We tested the reproducibility of robotic experiments, the iterative centrifugation and rinsing, and the automated ICP-MS measurements, which produced ion concentration measurements within a reasonable standard deviation of 4.0–7.6% and 2 outlier values from the batch of 96 samples.

Within a single iteration of these 96 experiments, 99.4% pure magnesium hydroxide salt was recovered at 86% yield upon addition of 3.4 equivalent of NaOH (Fig. 2a). Lower

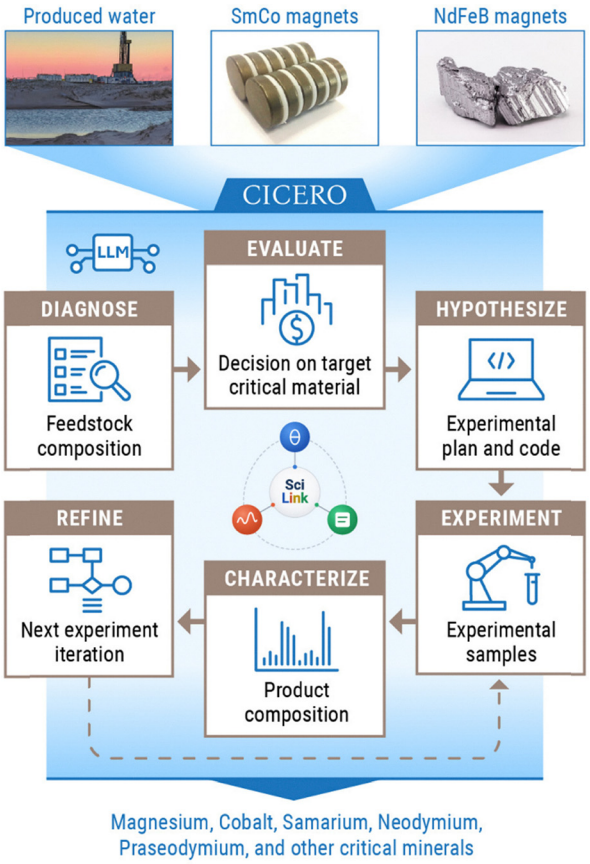


Fig. 1 Description of CICERO workflow, powered by multi-agentic SciLink platform. CICERO takes as input a real-world feedstock composition, and provides as output an optimized separation method to recover critical materials *via* selective precipitation.

[NaOH] also produced high-purity magnesium hydroxide, but with sub-optimal yields. Conversely, higher equivalents and concentrations of NaOH resulted in the inclusion of the competing alkaline earth metals, eventually forming precipitate mixtures that are 41.4% Mg, 56.5% Ca, 2.0% Sr, and 0.007% Ba calculated as mol % relative to all the metals in the product, reflecting the composition of the metal ions in the feedstock (Fig. 2b). To demonstrate the scalability from microplates to the bench scale, we performed bulk precipitation experiments using a thousand-fold larger volumes, reporting similar purity and yield of the magnesium product (details in SI, Fig. S2). Additional experiments in continuously stirred tank reactors at pilot scale and beyond would be needed to validate these results in the field. Moreover, the CICERO-led approach achieves direct magnesium recovery at yields and purity rates that compare with other precipitation-based methods, but without the need for removing calcium as a pre-treatment step.[26]

For the second feedstock, leached commercial SmCo magnets, the SciLink's Planning Agent identified both samarium and cobalt as valuable candidates for extraction, with samarium being the primary target. Unlike the produced water feedstock, here the recommended hypothesis involved two reagents, specifically adding both sodium bicarbonate and sodium hydroxide in various amounts. The underlying hypothesis was using a pH of

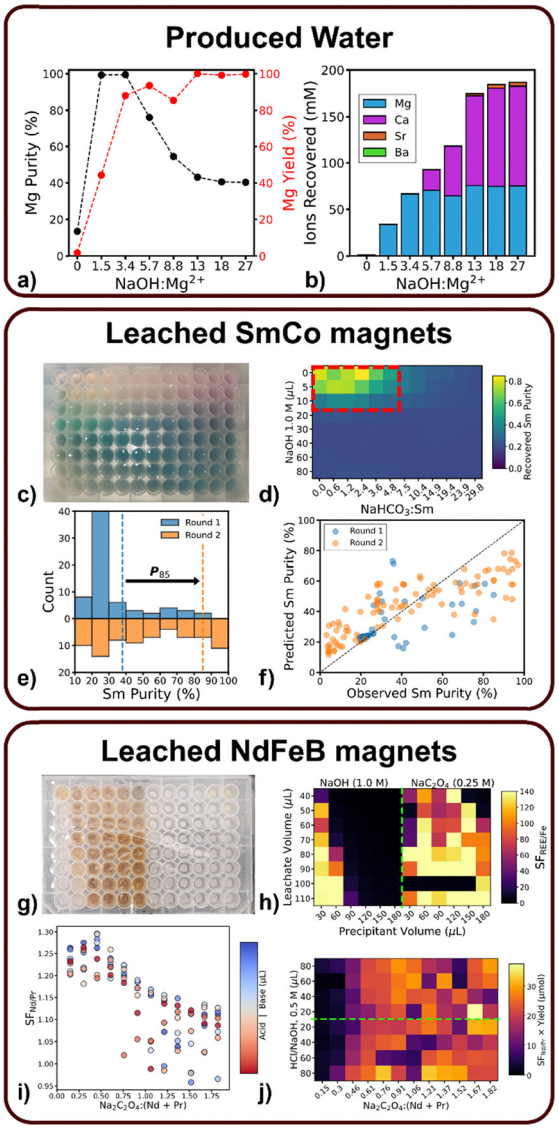


Fig. 2 Experimental results of critical materials separations from three real-world feedstocks, produced water, leached SmCo magnets, and leached NdFeB magnets. (a) Purity in mol% (black) and corresponding yield (red) of Mg recovered as precipitate *via* addition of NaOH solution to produced water. (b) Recovered Mg, Ca, Sr, and Ba from produced water as precipitates digested for ICP-MS measurements. (c) Image of plate containing SmCo feedstock after Round 1 of precipitation with $NaHCO_3$ and NaOH. Sm salts are colorless to faint yellow whereas Co salt precipitates range from pink to green and blue depending on crystalline phase of the hydroxide formed. (d) Purity of Sm salt recovered from the SmCo feedstock after the first round of precipitation experiments (where a value of 1 represents 100% purity relative to other metals). (e) Histogram of purities recovered from the first (blue) and second (orange) rounds of 96 precipitation experiments of SmCo feedstock. Dashed lines correspond to the 85th percentile of recovery for corresponding rounds. (f) Parity plot of the Bayesian Optimization machine learning on Sm% purity as a function of added $NaHCO_3$ and NaOH after the two rounds of data collection. (g) Image of a plate containing NdFeB feedstock after precipitation with NaOH (left) and $Na_2C_2O_4$ (right). REE precipitates are pale to colorless while Fe precipitates are orange. (h). Separation factor of Fe relative to combined REE total after Round 1 of precipitation of the NdFeB feedstock. (i) The separation factor of Nd from Pr after Round 2 of precipitation of NdFeB stock after Fe removal, as a function of added $Na_2C_2O_4$ and NaOH (blue) and HCl (red). (j) The separation factor from Figure multiplied by the yield.

approximately 6.5 to alter the speciation of cobalt complexes, while selectively precipitating samarium using the carbonate reagent. Such carbonate selectivity towards rare earth elements notably has literature precedent.[27] The SciLink's Planning Agent converted this hypothesis into a range of 96 experimental conditions and generated the corresponding Opentrons python script (SI). The parameter space spanned a range of $[NaHCO_3]$ = 5–200 mM and [NaOH] = 12.5–250 mM. The experiments resulted in products with a variety of colors; pink at high $[NaHCO_3]$ and low [NaOH], light brown at the opposite extreme of high [NaOH] and low $[(NH_4)_2CO_3]$, and a gradient of bluish green at most other conditions (Fig. 2c). These variations likely correspond to cobalt (oxy)hydroxide polymorphs formed at different solution conditions and were correlated with changes in the UV-vis spectra (Fig. S4). ICP-MS analysis showed reasonable separation after one round, achieving an 85.0% Sm salt product from an initially 21.9% Sm feedstock relative to the total metal ion composition, corresponding to a separation factor of 22.8, defined as $SF_{Sm/Co} = \frac{n_{Sm,f}/n_{Sm,i}}{n_{Co,f}/n_{Co,i}}$. (Fig. 2d). To further improve this outcome, we performed a second batch of experiments employing Sci-Link's built-in Bayesian optimization agent in conjunction with the planning agent. The second-round experimental plan and corresponding code suggested the design space of [NaOH] = 0.27–0.675 mM and $[NaHCO_3]$ = 20–80 mM, resulting in an increased purity of Sm product with a shift in the 85$^{th}$ percentile from ~40% in the first batch of experiments to 89%, at 92 > 99% yield (Fig. 2e and f), with noticeably improved confidence intervals in the high purity region and results that outcompete precipitation-based separation methods by few percentage points.[28]

The third feedstock that we investigated was the NdFeB magnet leachate, prompting yet a different approach from CICERO, which opted to explore two precipitating reagents separately, namely sodium hydroxide and sodium oxalate. The Planning Agent identified neodymium (23 800 ppm) and praseodymium (5900 ppm) as the primary targets, as well as the lower concentrations of dysprosium (27 ppm), terbium (16 ppm), and cobalt (470 ppm) as potential co-products, and attempted to bypass the substantial amounts of iron (62 000 ppm) and boron (920 ppm) in the feedstock. Our analysis shows that the oxalate was highly selective in recovering >99% pure rare earth element salts including both neodymium and praseodymium salts at 93% yield, reaching separation factors of approximately 200 relative to iron across a wide concentration regime (Fig. 2g and h). By comparison, hydroxide achieved reasonable separation factors at low precipitant concentrations, likely forming double sulfate salts as the magnet was leached in sulfuric acid,[29] but the yield was only 26%. At higher concentration, hydroxide precipitates incorporated iron in large amounts significantly compromising the purity. Since the removal of iron was achieved in the first round of oxalate experiments, no further Bayesian optimization was required. Instead, we opted to prompt the Planning Agent in the second round to improve the separation factor among the two rare earth elements, Nd and Pr. The agent suggested the continued use of oxalate with pH adjustment as a second tuning parameter. A single round of this strategy yielded a Nd/Pr

separation factor of up to 1.35 with a strong dependence on oxalate concentration and weak dependence on pH (Fig. 2i and j). This separation factor is competitive for rare earth element separations in a single round, especially since Nd anPr are adjacent elements in the periodic table and notoriously challenging to separate. The conditions were explored nearly autonomously using the workflow described earlier, over a period of approximately three days.

In conclusion, we have demonstrated a multi-agentic workflow that deploys a series of agents and automated instruments to achieve the separation and recovery of critical materials from a variety of complex feedstocks. Beyond the domain application, this work also contributed to the development of two new SciLink agents, the planning agent and the Bayesian optimization agent, for hypothesis generation, experimental planning, and closed-loop optimization. In the current workflow, the role of ''humans-in-the-loop'' is restricted to two categories of actions. The first is trivial tasks, such as prompting agents with preliminary data, clicking on instrument software buttons to conduct experiments, and moving the sample plates between the instruments. These tasks do not have significant bearing on the scientific nature of the workflow and current efforts to fully automate them are in progress. Such fully automated capabilities comparable to recently published work would be advantageous in accelerating throughout and iterative feedback.[30–32] We are in the process of developing a full digital twin of the laboratory described here, connected with robotic arms, that would facilitate such experimental campaigns. The other category of human intervention involves supervising and validating the output from key science agents that perform the hypothesis-generation, experiment generation, and data processing. This consideration is even more relevant when applying the workflow to problems that push the boundaries of the literature corpus, such as materials discovery, rather than the optimization of relatively simple chemistries based on speciation and precipitation as described here. In the current work, intervention was kept to a bare minimum to better reflect SciLink's capabilities, while checking that the agent output is rational; these interventions included correcting CICERO on the maximum concentration of sodium oxalate, adding a trash bin location to the Opentrons script, and suggesting working volumes be kept above 5 μL for precision. Some examples of human input are provided in Fig. S8 and the corresponding SI discussion.

Building on the current version of the workflow, we anticipate that the ''Evaluate'' and ''Hypothesize'' steps could be incorporated in the iterative loop after the results are analysed, such that the agent can decide to target a different critical material in the feedstock or propose an alternative hypothesis for extracting the target critical material based on how the experimentation is progressing. Firstly, this iteration mitigates the risk of an initially bad hypothesis producing a narrow parameter space that is sub-optimal for separations, yet challenging for the agents to backtrack and reason out of. Moreover, this could help identify opportunities for recovering additional critical materials that were not initially designated as the targets, but whose recovery becomes viable as new experimental results are processed. Substantial research efforts continue to focus on achieving and optimizing critical materials separations *via* selective precipitation; CICERO demonstrates how the efficiency of this research can be massively enhanced using agentic workflows and automated experiments. Furthermore, while the current implementation uses Gemini-2.5-Pro, the workflow is architecturally model-agnostic, as SciLink agents interface with LLMs through a standardized API layer. We therefore expect CICERO to benefit and improve naturally from ongoing advances in both proprietary and open-source LLMs.

Lastly, CICERO's ability to rapidly pivot between fundamentally different separation strategies for each feedstock demonstrates the adaptability that makes agentic approaches particularly valuable when processing diverse, real-world feedstocks. For the produced water feedstock, near-optimal recovery of magnesium was achieved from the first batch of experiments without the need for additional experimental iterations. For the SmCo feedstock, the agent opted for an additional round of Bayesian optimization to improve the separation of the two primary elements in that mixture. The NdFeB prompted yet a different response, with the agent finding better value in separating individual REE elements from each other, rather than further optimizing the separation of the REE group from other impurities. This varied approach is an indicator of an emerging trend in AI-drive science, where LLMs coupled to Bayesian optimization can find more versatile experimental strategies, beyond focusing on a single objective that is pre-defined by the human user.[33] Furthermore, the data generated by CICERO – with full experimental provenance from hypothesis to execution to analysis and decision-making – creates a reproducible knowledge base that can inform future separations research. As global demand for critical materials intensifies and feedstock complexity increases, agentic workflows like CICERO offer a scalable path toward rapidly developing feedstock-specific, economically viable separation processes that would be impractical to optimize through traditional approaches alone.

## Conflicts of interest

There are no conflicts to declare.

## Data availability

The data supporting this article have been appended as part of the supplementary information (SI), including JSON files, instrument characterization files, scripts for Opentrons operation, and agent prompts and responses. More on SciLink can be found on https://github.com/ziatdinovmax/scilink. See DOI: https://doi.org/10.1039/d6mh00475j.

## Acknowledgements

Feedstock acquisition and separation experiments were supported by the Non-Equilibrium Transport Driven Separations (NETS) initiative, under the Laboratory Directed Research and Development (LDRD) Program at Pacific Northwest National

Laboratory (PNNL). Data processing and management was supported by the Adaptive Tunability for Synthesis and Control *via* Autonomous Learning on Edge (ATSCALE) initiative, under the LDRD program at PNNL. The early development of SciLink was supported by seed LDRD funding from the Physical and Computational Sciences Directorate at PNNL. The setup of the automated instruments was supported by the Foundational Autonomy Investment (FAI), under the LDRD program at PNNL. PNNL is a multi-program national laboratory operated for the U.S. Department of Energy (DOE) by Battelle Memorial Institute under Contract No. DE-AC05-76RL01830. The graphic in Fig. 1 was developed by Nathan Johnson.

## References

1 D. S. Sholl and R. P. Lively, Seven chemical separations to change the world, *Nature*, 2016, **532**, 435–437.
2 *National Academies of Sciences, Engineering, and Medicine: A Research Agenda for Transforming Separation Science*, The National Academies Press, Washington, DC, 2019.
3 Y. Yang, A. Walton, R. Sheridan, K. Güth, R. Gauß, O. Gutfleisch, M. Buchert, B.-M. Steenari, T. Van Gerven and P. T. Jones, *et al.*, REE Recovery from End-of-Life NdFeB Permanent Magnet Scrap: A Critical Review, *J. Sustainable Metall.*, 2016, **3**(1), 122–149, DOI: 10.1007/s40831-016-0090-4.
4 Q. Wang, E. Nakouzi, E. A. Ryan and C. V. Subban, Flow-Assisted Selective Mineral Extraction from Seawater, *Environ. Sci. Technol. Lett.*, 2022, **9**, 645–659.
5 Q. Wang, Y. Fu, E. A. Miller, D. Song, P. J. Brahana, A. Ritchhart, Z. Xu, G. E. Johnson, B. Bharti and M. L. Sushko, *et al.*, Selective Recovery of Critical Minerals from Simulated Electronic Wastes *via* Reaction-Diffusion Coupling, *ChemSusChem*, 2025, e202402372, DOI: 10.1002/cssc.202402372.
6 P. Butreddy, S. T. Mergelsberg, J. N. Jocz, D. Li, V. Prabhakaran, A. J. Ritchhart, C. V. Subban, J. Kellar, S. R. Beeler and S. W. Keenan, *et al.*, Selective dissolution and re-precipitation by pH cycling enables recovery of manganese from surface nodules, *RSC Sustainability*, 2025, **3**, 983–994, DOI: 10.1039/d4su00444b.
7 Y. Shekarian, M. Rezaee and S. Pisupati, Green chemical precipitation of manganese, cobalt, and nickel from acid mine drainage using ozone: mechanism and chemical kinetics, *React. Chem. Eng.*, 2025, **10**(10), 2398–2411, DOI: 10.1039/d5re00222b.
8 X. Zhang, L. Zeng, Y. Wang, J. Tian, J. Wang, W. Sun, H. Han and Y. Yang, Selective separation of metals from wastewater using sulfide precipitation: A critical review in agents, operational factors and particle aggregation, *J. Environ. Manage.*, 2023, **344**, 118462, DOI: 10.1016/j.jenvman.2023.118462 From NLM Publisher.
9 Q. Li, Z. Xiao and W. Zhang, Sulfide precipitation characteristics of Mn, Ni, Co, and Zn in the presence of contaminant metal ions, *Miner. Eng.*, 2024, 215, DOI: 10.1016/j.mineng.2024.108814.
10 P. Sun, D. Jing, D. B. Durham, B. Sapkota, Y. Chen and J. L. Anderson, Ion-Specific Precipitation of Extractants Enables Rare-Earth Separation and Wastewater Remediation from Solvent Extraction of Critical Elements, *Environ. Sci. Technol.*, 2026, **60**, 6239–6248, DOI: 10.1021/acs.est.5c16966From NLM Publisher.
11 L. Xu, B. Zhao, X. Zhang, W. Liu, D. Rau, D. Wu, W. Zhang, C. Liu, Z. Liu and S. Lin, Membrane and electrochemical separations for direct lithium extraction, *Nat. Chem. Eng.*, 2025, **2**, 551–567.
12 O. A. Kazi, W. Chen, J. G. Eatman, F. Gao, Y. Liu, Y. Wang, Z. Xia and S. B. Darling, Material Design Strategies for Recovery of Critical Resources from Water, *Adv. Mater.*, 2023, **35**(36), e2300913, DOI: 10.1002/adma.202300913 From NLM PubMed-not-MEDLINE.
13 M. Kaya, An overview of NdFeB magnets recycling technologies, *Curr. Opin. Green Sustainable Chem.*, 2024, **46**, 100884.
14 S. E. Can Sener, V. M. Thomas, D. E. Hogan, R. M. Maier, M. Carbajales-Dale, M. D. Barton, T. Karanfil, J. C. Crittenden and G. L. Amy, Recovery of Critical Metals from Aqueous Sources, *ACS Sustainable Chem. Eng.*, 2021, **9**(35), 11616–11634, DOI: 10.1021/acssuschemeng.1c03005 From NLM PubMed-not-MEDLINE.
15 A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser and I. Polosukhin Attention is All you Need. In NeurIPS, 2017.
16 A. Radford, J. Wu, R. Child, D. Luan, D. Amodei and I. Sutskever Language Models are Unsupervised Multitask Learners. Semantic Scholar 2019, CorpusID:160025533.
17 P. F. Christiano, J. Leike, T. B. Brown, M. Martic, S. Legg and D. Amodei Deep reinforcement learning from human preferences. In NeurIPS, 2017.
18 J. Devlin, M.-W. Chang, K. Lee and K. Toutanova BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. NAACL-HLT, 2019, pp 4171–4186.
19 A. Radford and K. Narasimhan Improving Language Understanding by Generative Pre-Training. Semantic Scholar 2018, CorpusID:49313245.
20 T. M. Brown, B. Ryder, N. Subbiah, M. Kaplan, J. D. Dhariwal, P. Neelakantan, A. Shyam, P. Sastry, G. Askell, A. Agarwal, A. Herbert-Voss, A. Krueger, G. Henighan, T. Child, R. Ramesh, A. Ziegler, D. Wu, J. Winter, C. Hesse, C. Chen, M. Sigler, E. Litwin, M. Gray, S. Chess, B. Clark, J. Berner, C. McCandlish, S. Radford, A. Sutskever and I. Amodei, D. Language Models are Few-Shot Learners. In NeurIPS, 2020.
21 J. Wei, X. Wang, D. Schuurmans, M. Bosma, B. Ichter, F. Xia, E. H. Chi, Q. V. Le and D. Zhou Chain-of-thought prompting elicits reasoning in large language models. *In NeurIPS*, 2022.
22 L. Ouyang, J. Wu, X. Jiang, D. Almeida, C. Wainwright, P. Mishkin, C. Zhang, S. Agarwal, K. Slama, A. Ray, *et al.* Training language models to follow instructions with human feedback. In NeurIPS, 2022.
23 L. Yao, S. Samantray, A. Ghosh, K. Roccapriore, L. Kovarik, S. Allec and M. Ziatdinov *Operationalizing Serendipity: Multi-Agent AI Workflows for Enhanced Materials Characterization with Theory-in-the-Loop*, 2025, preprint, arXiv:2508.06569.
24 B. Shahriari, K. Swersky, Z. Wang, R. P. Adams and N. de Freitas, Taking the Human Out of the Loop: A Review of Bayesian Optimization, *Proc. IEEE*, 2015, **104**, 148–175.
25 al., A. e. *Gemini: A Family of Highly Capable Multimodal Models*, 2023, preprint, arXiv:2312.11805.
26 L. Hu, J. Yu, H. Luo, H. Wang, P. Xu and Y. Zhang, Simultaneous recovery of ammonium, potassium and

 

magnesium from produced water by struvite precipitation, *Chem. Eng. J.*, 2020, 382, DOI: 10.1016/j.cej.2019.123001.

27 X. Ren, L. Wei, A. Sun, B. Xu, W. Jiang, B. Yang and F. Wang, Recycling of Sm-Co permanent magnet waste: Perspectives and recent advances, *J. Rare Earths*, 2025, DOI: 10.1016/j.jre.2025.09.038.

28 E. Uysal, E. Emil-Kaya, H. N. Dursun, M. Papakci, S. Gürmen and B. Friedrich, Recovery of Samarium from Waste SmCo Magnets *via* Selective Precipitation with Ammonium Bicarbonate: Optimization of Process Efficiency, *Metals*, 2024, **14**(12), 1363, DOI: 10.3390/met14121363.

29 J. W. Lyman and G. R. Palmer, Recycling of Rare Earths and Iron from NdFeB Magnet Scrap, *High Temp. Mater. Processes*, 1993, **11**, 175–187.

30 N. Yoshikawa, Y. Asano, D. N. Futaba, K. Harada, T. Hitosugi, G. N. Kanda, S. Matsuda, Y. Nagata, K. Nagato and M. Naito, *et al.*, Self-driving laboratories in Japan, *Digital Discovery*, 2025, **4**(6), 1384–1403, DOI: 10.1039/d4dd00387j.

31 Y. Xie, K. Sattari, C. Zhang and J. Lin, Toward autonomous laboratories: Convergence of artificial intelligence and experimental automation, *Prog. Mater. Sci.*, 2023, 132, DOI: 10.1016/j.pmatsci.2022.101043.

32 N. Ishizuki, R. Shimizu and T. Hitosugi, Autonomous experimental systems in materials science, *Sci. Technol. Adv. Mater.:-Methods*, 2023, 2197519, DOI: 10.1080/27660400.2023.2197519.

33 A. Ritchhart, S. Allec, H. Job, P. Butreddy, M. Ziatdinov and E. Nakouzi, Toward autonomous materials synthesis *via* reaction–diffusion coupling, *APL Mach. Learn.*, 2025, **3**, 036109.