
Closed-loop, machine learning–driven optimization of reactor yields in reactive carbon electrolyzers

Abhishek Soni¹, Siwei Ma¹, Andrew Jewlal¹, Mehrdad Mokhtari¹, Ribwar Ahmadi¹,
Daniel Lin¹, Karry Ocean¹, Giuseppe V. Crescenzo¹, Kevan Dettelbach¹,
Nada Nasr¹, Curtis Berlinguette¹

¹The University of British Columbia, Vancouver, BC, Canada, V6T 1Z1*

Abstract

Reactive carbon capture combines CO₂ capture and conversion in a single system. Reactive carbon electrolyzers receive a liquid eluent from a CO₂ capture unit containing a sorbent that has captured CO₂. This electrolyzer releases CO₂ electrochemically and converts it into a value-added product like CO. The effectiveness of this system depends on high CO₂ utilization and high product formation rates. We define their product as “reactor yield.” Here, we used a closed-loop, automated workflow with Bayesian optimization to maximize reactor yield in an electrolyzer operating with alkaline CO₂ capture solutions. We explored a six-dimensional parameter space and found that a bicarbonate concentration of 1.5 M and carbonate concentration of 0.75 M achieved the highest reactor yield (44 mA cm^{−2}). Interestingly, this optimum occurred at non-maximum values of CO partial current density (54 vs. 87 mA cm^{−2}) and CO₂ utilization (81% vs. 100%), highlighting the need for joint optimization of both factors.

1 Introduction

Carbon dioxide (CO₂) capture and conversion technologies transform captured CO₂ into higher-value carbon products (e.g., CO, C₂H₄). [1, 2, 3, 4, 5, 6] Current capture and conversion technologies require pressurization and thermal steps to purify and release captured CO₂ before use. [7, 8, 9] Reactive carbon capture avoids these steps by using electrolysis to capture and convert CO₂ in a closed-loop. [10]

Kim et al previously demonstrated reactive carbon capture by using a caustic aqueous capture solution (e.g., OH[−]) to capture CO₂ from air as a (bi)carbonate-enriched solution. [11] They fed this solution into the electrolyzer. The (bi)carbonate ions in solution (HCO₃[−] and CO₃^{2−}) then react with electrolytically generated protons in the cathodic compartment to produce CO₂ in situ (*i*-CO₂). Next, *i*-CO₂ is electrochemically reduced to higher-value products, such as CO, CH₄, [12, 13] or C₂H₄. [14, 15, 16] The caustic liquid exiting the electrolyzer is recirculated into a capture unit to capture more CO₂. [11, 17] This closed-loop workflow requires less infrastructure and enables the use of liquid feedstocks that simplify engineering, operation, and impurity management. [10]

We evaluate the performance of a reactive carbon electrolyzer by tracking the partial current density toward CO (J_{CO} ; equation (1)), and CO₂ utilization (U_{CO_2} ; equation (2)). J_{CO} is a measure of CO formation rate, while U_{CO_2} indicates how efficiently the CO₂ is consumed:

$$J_{CO} \text{ (mA cm}^{-2}\text{)} = J_{\text{total}} \times FE_{CO} \quad (1)$$

*39th Conference on Neural Information Processing Systems (NeurIPS 2025) Workshop: AI4Mat.




Figure 1: Closed-loop platform for optimizing reactor yield in a reactive carbon electrolyzer. (A) AdaCarbon automates the optimization of reactor yield by adjusting six operating parameters in an automated test cell (ATC). The ATC collects data, which is processed and sent to a Bayesian optimization algorithm that selects the next experiment. This loop continues until reactor yield stops improving. (B) Optimization landscape showing reactor yield as a function of partial current density for CO (J_{CO}) and CO₂ utilization (U_{CO_2}). Reactor yield peaks at intermediate values of both metrics, reflecting a trade-off between J_{CO} and U_{CO_2} .

$$U_{CO_2}(\%) = \frac{n_{CO_2RR}}{n_{CO_2RR} + n_{CO_2}} \times 100\% \quad (2)$$

where FE_{CO} is the Faradaic efficiency toward CO, n_{CO_2RR} is the molar amount of CO₂ reduction products exiting the reactor, and n_{CO_2} is the molar amount of CO₂ exiting the reactor. These two performance metrics trade off, where higher J_{CO} reduces U_{CO_2} due to excess CO₂ generation or poor mass transport. This trade-off complicates efforts to improve electrolyzer performance. To resolve this trade-off, we define the figure of merit “reactor yield”. The reactor yield is the product of J_{CO} and U_{CO_2} (equation (3)):

$$Reactor\ yield\ (\text{mA cm}^{-2}) = J_{CO} \times U_{CO_2} \quad (3)$$

To maximize reactor yield, the CO production rate and CO₂ utilization must be balanced. Each new operating parameter expands the parameter space exponentially, making manual optimization impractical.[18, 19] We used our self-driving laboratory, “AdaCarbon”, [20, 21] to navigate this parameter space and optimize reactor yield for reactive carbon electrolysis (Fig. 1A). AdaCarbon automates electrode fabrication, electrolysis, and data acquisition. Coupled with Bayesian optimization, AdaCarbon performs closed-loop experiments that iteratively improve reactor yield (Fig. 1B). We show that optimizing for reactor yield requires different conditions than optimizing for J_{CO} or U_{CO_2} alone. Our results provide a framework for tuning reactive carbon electrolyzers using flexible automation to optimize reactor yield.

2 Results

2.1 Parameter selection and optimization framework

We used AdaCarbon to optimize reactor yield by varying six experimental parameters: bicarbonate temperature, concentration, and flow rate, KOH concentration and flow rate, and total applied current density (Supplementary Table 1; Supplementary Figs. 1–3; and Methods section). Prior reports link some of these parameters individually to J_{CO} and U_{CO_2} . [22, 23, 24, 25]

Bicarbonate temperature (25–80 °C) affects the CO₂ solubility.[22] Bicarbonate concentration (0.5–3.05 M) controls the availability of *i*-CO₂ for electrochemical reduction.[23] Bicarbonate




Figure 2: Closed-loop optimization of reactor yield and J_{CO} using AdaCarbon. (A) AdaCarbon initialized the reactor yield optimization campaign with 12 random experiments, followed by 10 optimizer-guided experiments. Five researcher-defined (RD) experiments were then added to train the optimizer further. After 37 total runs, the reactor yield peaked at 38 mA cm^{-2} (star). More experiments afterwards did not improve reactor yield. (B) The partial current density for CO (J_{CO}) optimization campaign initialized with all 52 reactor yield experiments. After five more runs, the optimizer identified a maximum J_{CO} of 140 mA cm^{-2} . (C) Comparison of the two campaigns. Intermediate J_{CO} and high CO_2 utilization (U_{CO_2}) were required to optimize reactor yield. Contour lines show reactor yield as a function of J_{CO} and U_{CO_2} .

flow rate (30–200 mL min^{-1}) sets the residence time of $i\text{-CO}_2$ near the catalyst surface.[24] KOH concentration (0.05–3.00 M) modulates the cathodic microenvironment,[25] and KOH flow rate (6–30 mL min^{-1}) was included to probe similar effects. The total applied current density (50–350 mA cm^{-2}) shifts product selectivity (i.e., FE).[10] We chose the parameter ranges from manual experiments (Supplementary Fig. 4) and physical constraints (e.g. membrane stability).

Manual exploration of this six-dimensional parameter space is infeasible. Even with five values per parameter, the combinations exceed 15,000 experiments to test (Supplementary Fig. 5). We addressed this challenge using an automated AdaCarbon platform equipped with machine learning and Bayesian optimization (Supplementary Note 1; Methods section).[26] This enabled AdaCarbon to execute experiments, analyze results, and select the next experiments without human intervention.

We trained a surrogate model using Gaussian process regression[27] to predict reactor yield from the six input parameters. The optimizer utilized a q-noisy expected improvement acquisition function to select the next experiment.[28] This algorithm prioritized regions most likely to increase reactor yield and undersampled areas. By accounting for uncertainty in noisy electrolysis data, this algorithm improved accuracy and robustness (Supplementary Figs. 6–7; Supplementary Note 2).

2.2 Closed-loop reactor yield optimization campaign and champion conditions

We initialized the reactor yield optimization campaign with 12 experiments using randomly distributed parameter sets (Fig. 2A; Methods section). These experiments trained the Bayesian optimizer to

predict reactor yield across the six-dimensional parameter space. Before each run, a researcher loaded fresh components (bipolar membrane, nickel foam, and silver gas-diffusion electrodes) into the automated test cell (ATC; Fig. 1 and Supplementary Fig. 2). Electrolysis ran for 5 minutes and products were analyzed through mass spectrometry, then processed and used by the optimizer to inform the next iteration.

The optimizer proposed 10 additional experiments after the 12 initial runs, bringing the total to 22. When compared to initialization runs, these runs increased reactor yield by 4%—yielding a 33 mA cm^{-2} reactor yield (Fig. 2A). To probe the parameter space further and test hypotheses, we added 5 researcher-defined (RD) experiments (Fig. 2A; Supplementary Table 2; Supplementary Note 3), raising the total to 27. The optimizer then identified a peak reactor yield of 38 mA cm^{-2} only in 37 experiments (Fig. 2A; star)—over 3.5 times higher than the best reactor yield from prior spray-coated silver catalysts.[29] No further gains were observed after 52 total experiments, indicating convergence (Supplementary Fig. 8).

The champion conditions that yielded 38 mA cm^{-2} balanced U_{CO_2} (52%) and J_{CO} (73.5 mA cm^{-2}). These values did not coincide with the maxima for either metric— J_{CO} peaked at 92 mA cm^{-2} and U_{CO_2} at 56% (Supplementary Table 3)—reinforcing the trade-off between them. We confirmed this trade-off during a separate sequential optimization campaign targeting J_{CO} (Fig. 2B). At the peak J_{CO} of 140 mA cm^{-2} , achieved after only five optimization runs, U_{CO_2} dropped to 13% (Fig. 2C). As a result, the reactor yield decreased to 18 mA cm^{-2} (Supplementary Fig. 9).




Figure 3: **SHAP analysis shows how input parameters affect reactor yield and J_{CO} .** (A) In the reactor yield campaign, bicarbonate concentration and current density were most influential. These parameters control $i\text{-CO}_2$ formation and J_{CO} . (B) In the J_{CO} campaign, the same variables dominated, but the optimizer selected higher values, increasing J_{CO} and decreasing U_{CO_2} . Positive SHAP values indicate direct correlation with the target metric; negative values indicate inverse correlation. Each point represents one experiment.

2.3 Reactor yield and J_{CO} sensitivity to input parameters

We used SHAP (SHapley Additive exPlanations) analysis[30, 31] to understand which parameters most influence reactor yield and J_{CO} . SHAP values quantify the individual impact of each parameter, where the sign (+/-) indicates correlation direction and the magnitude indicates strength (see Supplementary Note 4 for details).

In the reactor yield campaign, bicarbonate concentration and applied current density were most influential (Fig. 3A) because they directly affect in situ CO_2 ($i\text{-CO}_2$) formation and J_{CO} . [29] The optimizer selected low bicarbonate concentrations and moderate current densities (Supplementary Fig. 10A–B), challenging the assumption that higher bicarbonate concentrations always improve performance.[16, 17] Lower bicarbonate concentrations improved U_{CO_2} while moderate current densities balanced product selectivity and, therefore, J_{CO} . By balancing this trade-off, the optimizer improved reactor yield.

In the J_{CO} campaign, higher bicarbonate concentration and total applied current density again dominated (Fig. 3B), but the optimizer favoured higher values (Supplementary Fig. 11A–B). These conditions increased $i\text{-CO}_2$ and J_{CO} but sharply reduced U_{CO_2} .

In both campaigns, moderate to high bicarbonate flow rates improved reactor yield (Supplementary Figs. 10C and 11C). In the J_{CO} campaign, the optimizer mitigated the typical drop in FE_{CO} at high current densities (Supplementary Fig. 12; orange points). We attribute this to improved $i\text{-CO}_2$ bubble management and shorter residence times, enhancing $i\text{-CO}_2$ mass transport and reaction rates.[32]

The KOH concentration, KOH flow rate, and bicarbonate temperature had limited influence on reactor yield and J_{CO} . Although KOH conditions influence ion transport and catalyst behaviour,[25] their influence was small compared to bicarbonate-related variables and current density. Temperature effects on CO_2 solubility were also minimal.

2.4 Demonstration of optimized reactor yield under realistic capture conditions

We tested reactor yield under realistic alkaline CO_2 capture conditions to assess the relevance of the optimized conditions under practical carbon capture scenarios.[33] The electrolyte contained varying concentrations of bicarbonate and carbonate, with a fixed 3 M K^+ concentration (Supplementary Note 5). We applied the champion experimental conditions from the reactor yield optimization campaign (Supplementary Table 3) and varied the bicarbonate:carbonate ratio (Supplementary Table 4). The reactor yield peaked at 44 mA cm^{-2} with 1.5 M KHCO_3 and 0.75 M K_2CO_3 (Fig. 4, top panel).




Figure 4: **Reactor yield optimization in a realistic reactive carbon capture electrolyte.** Top: Reactor yield plotted as a function of bicarbonate concentration in a fixed 3 M K^+ electrolyte. Carbonate concentration was adjusted to maintain total K^+ (Supplementary Table 4). Reactor yield peaked at 44 mA cm^{-2} with 1.5 M KHCO_3 and 0.75 M K_2CO_3 . Bottom: Corresponding J_{CO} and U_{CO_2} values. Pure carbonate (0 M bicarbonate) resulted in 100% U_{CO_2} but low J_{CO} (17 mA cm^{-2}). Pure bicarbonate (3 M) gave high J_{CO} (87 mA cm^{-2}) but low U_{CO_2} (27%). The maximum reactor yield was observed with intermediate values of 54 mA cm^{-2} J_{CO} and 81% U_{CO_2} , demonstrating the trade-off.

These experiments further confirmed the trade-off between J_{CO} and U_{CO_2} (Fig. 4, bottom panel). With pure carbonate (0 M bicarbonate), U_{CO_2} reached 100%, while J_{CO} dropped to 17 mA cm^{-2} . With pure bicarbonate (3 M), J_{CO} rose to 87 mA cm^{-2} but U_{CO_2} fell to 27%. The optimal reactor yield of 44 mA cm^{-2} was achieved with intermediate concentrations of 1.5 M KHCO_3 and 0.75 M

K₂CO₃. These conditions balanced J_{CO} (54 mA cm⁻²) and U_{CO_2} (81%), maximizing reactor yield. These results demonstrate that reactor yield remains a robust metric for optimizing reactive carbon capture.

3 Conclusion

We have shown that optimizing the reactor yield of reactive carbon electrolyzers identifies reaction conditions that balance the trade-off between product formation rates (J_{CO}) and CO₂ utilization. This contrasts with traditional optimization targets like J_{CO} , which ignore efficiency losses from unreacted CO₂ and upstream separation costs.

By using an automated AdaCarbon platform equipped with machine learning and Bayesian optimization, we navigated a six-variable parameter space to converge on conditions that maximize reactor yield in fewer than 40 experiments. We validated these results in a realistic mixed bicarbonate-carbonate electrolyte, where trade-offs between J_{CO} and CO₂ utilization emerged. Optimal reactor yield was achieved at intermediate J_{CO} and CO₂ utilization—not at their extremes.

This work demonstrates that autonomous optimization of reactor yield offers a scalable path to unify CO₂ capture and conversion, opening new directions in carbon-neutral chemical production. Future work should automate catalyst development to further improve reactor yield by enhancing *i*-CO₂ conversion. Self-driving labs like AdaCarbon will be essential to accelerate the next-generation of catalyst discovery.

4 Data availability

The raw and processed data generated by the self-driving laboratory in this study are available in the supplementary information and at <https://github.com/berlinguette/ada>. All other data related to this paper are available from the corresponding author upon request.

5 Code availability

All code used in this study was based on open-source Python packages listed in the supplementary information and are available at <https://github.com/berlinguette/ada>.

References

- [1] Shaoyun Hao, Ahmad Elgazzar, Shou-Kun Zhang, Tae-Ung Wi, Feng-Yang Chen, Yuge Feng, Peng Zhu, and Haotian Wang. Acid-humidified CO₂ gas input for stable electrochemical CO₂ reduction reaction. *Science*, 388(6752):eadr3834, 2025.
- [2] Uzoma O. Nwabara, Emiliana R. Cofell, Sumit Verma, Emanuela Negro, and Paul J. A. Kenis. Durable cathodes and electrolyzers for the efficient aqueous electrochemical reduction of CO₂. *ChemSusChem*, 13(5):855–875, 2020.
- [3] A. Soni, X. Lu, C. Zhou, S. Singh, and C. P. Berlinguette. Thermoelectric CO₂RR electrolysis. *Device*, 2:100603, 2024.
- [4] Y. Wu, Z. Jiang, X. Lu, Y. Liang, and H. Wang. Domino electroreduction of CO₂ to methanol on a molecular catalyst. *Nature*, 575:639–642, 2019.
- [5] B. S. Crandall et al. Kilowatt-scale tandem CO₂ electrolysis for enhanced acetate and ethylene production. *Nat. Chem. Eng.*, 1:421–429, 2024.
- [6] D. M. Weekes, D. A. Salvatore, A. Reyes, A. Huang, and C. P. Berlinguette. Electrolytic CO₂ reduction in a flow cell. *Acc. Chem. Res.*, 51:910–918, 2018.
- [7] E. S. Sanz-Pérez, C. R. Murdock, S. A. Didas, and C. W. Jones. Direct capture of CO₂ from ambient air. *Chem. Rev.*, 116:11840–11876, 2016.

- [8] D. W. Keith, G. Holmes, D. St. Angelo, and K. Heidel. A process for capturing CO₂ from the atmosphere. *Joule*, 2:1573–1594, 2018.
- [9] S. Deutz and A. Bardow. Life-cycle assessment of an industrial direct air capture process based on temperature–vacuum swing adsorption. *Nat. Energy*, 6:203–213, 2021.
- [10] D. J. D. Pimlott, Y. Kim, and C. P. Berlinguette. Reactive carbon capture enables CO₂ electrolysis with liquid feedstocks. *Acc. Chem. Res.*, 57:1007–1018, 2024.
- [11] Y. Kim et al. Integrated CO₂ capture and conversion to form syngas. *Joule*, 8:3106–3125, 2024.
- [12] C. A. Obasanjo et al. High-rate and selective conversion of CO₂ from aqueous solutions to hydrocarbons. *Nat. Commun.*, 14:3176, 2023.
- [13] E. W. Lees et al. Electrolytic methane production from reactive carbon solutions. *ACS Energy Lett.*, 7:1712–1718, 2022.
- [14] G. Lee et al. CO₂ electroreduction to multi-carbon products from carbonate capture liquid. *Joule*, 7:1277–1288, 2023.
- [15] S. Ma et al. Cathode surface pH modulates multicarbon product selectivity during the electrochemical conversion of CO₂ capture solutions. *ACS Energy Lett.*, 9:2326–2332, 2024.
- [16] A. Venkataraman et al. Process and techno-economic analyses of ethylene production by electrochemical reduction of aqueous alkaline carbonates. *Nat. Chem. Eng.*, 2024.
- [17] H. M. Almajed et al. Closing the loop: Unexamined performance trade-offs of integrating direct air capture with (bi)carbonate electrolysis. *ACS Energy Lett.*, 9:2472–2483, 2024.
- [18] R. B. Canty et al. Science acceleration and accessibility with self-driving labs. *Nat. Commun.*, 16:3856, 2025.
- [19] B. P. MacLeod, F. G. L. Parlane, A. K. Brown, J. E. Hein, and C. P. Berlinguette. Flexible automation accelerates materials discovery. *Nat. Mater.*, 21:722–726, 2022.
- [20] I. Grega et al. A physics-based data-driven model for CO₂ gas diffusion electrodes to drive automated laboratories, arXiv, [cond-mat.mtrl-sci], 2025, doi:10.48550/arXiv.2502.06323.
- [21] A. Soni et al. Accelerated development of gas diffusion electrodes for CO₂ electrolyzers. ChemRxiv, 2024, doi:10.26434/chemrxiv-2024-znh0f.
- [22] Z. Zhang et al. Porous metal electrodes enable efficient electrolysis of carbon capture solutions. *Energy Environ. Sci.*, 15:705–713, 2022.
- [23] C. Larrea, D. Torres, J. R. Avilés-Moreno, and P. Ocón. Multi-parameter study of CO₂ electrochemical reduction from concentrated bicarbonate feed. *J. CO₂ Util.*, 57:101878, 2022.
- [24] O. Gutiérrez-Sánchez, B. de Mot, M. Bulut, D. Pant, and T. Breugelmanns. Engineering aspects for the design of a bicarbonate zero-gap flow electrolyzer for the conversion of CO₂ to formate. *ACS Appl. Mater. Interfaces*, 14:30760–30771, 2022.
- [25] K. Yang et al. Cation-driven increases of CO₂ utilization in a bipolar membrane electrode assembly for CO₂ electrolysis. *ACS Energy Lett.*, 6:4291–4298, 2021.
- [26] A. Hebbal, L. Brevault, M. Balesdent, E.-G. Talbi, and N. Melab. Bayesian optimization using deep gaussian processes, arXiv [stat.ML], 2019.
- [27] C. E. Rasmussen and C. K. I. Williams. *Gaussian Processes for Machine Learning*. MIT Press, 2005.
- [28] B. Letham, B. Karrer, G. Ottoni, and E. Bakshy. Constrained bayesian optimization with noisy experiments. *Bayesian Anal.*, 14:495–519, 2019.
- [29] T. Li et al. Electrolytic conversion of bicarbonate into CO in a flow cell. *Joule*, 3:1487–1497, 2019.

- [30] SHAP. Welcome to the SHAP documentation — SHAP latest documentation. <https://shap.readthedocs.io/en/latest/>.
- [31] A. Soni and G. N. Patey. Using machine learning with atomistic surface and local water features to predict heterogeneous ice nucleation. *J. Chem. Phys.*, 160:124501, 2024.
- [32] X. Lu et al. Visualization of CO₂ electrolysis using optical coherence tomography. *Nat. Chem.*, 16:979–987, 2024.
- [33] I. Sullivan et al. Coupling electrochemical CO₂ conversion with CO₂ capture. *Nat. Catal.*, 4:952–958, 2021.
- [34] C. C. Rupnow et al. A self-driving laboratory optimizes a scalable process for making functional coatings. *Cell Rep. Phys. Sci.*, 4:101411, 2023.
- [35] S. Ren et al. Molecular electrocatalysts can mediate fast, selective CO₂ reduction in a flow cell. *Science*, 365:367–369, 2019.
- [36] M. Balandat et al. Botorch: A framework for efficient monte-carlo bayesian optimization, arXiv [cs.LG], 2019.
- [37] J. R. Gardner, G. Pleiss, D. Bindel, K. Q. Weinberger, and A. G. Wilson. Gpytorch: Blackbox matrix-matrix gaussian process inference with gpu acceleration, 2018.

A Appendix

A.1 Methods

A.1.1 Reagents

Silver nanoparticles (25 nm, 99.5% trace metals basis) were obtained from Sigma Aldrich. Nickel foam (>99.99%) was sourced from MTI Co. Bipolar membranes (Fumasep FBM) were purchased from the Fuel Cell Store (USA) and stored in 1 M NaCl Solution for 24 h before use. KHCO_3 and KOH (99%) were obtained from Sigma-Aldrich. N_2 (99%), CO_2 (99%), CO (99%), and Ar (99.999%) gasses were supplied by Praxair Canada Inc. Freudenberg H23 carbon paper, and NafionTM Dispersion (D-2020 dispersion, 20% w/w in water and isopropyl alcohol) were purchased from the Fuel Cell Store (USA). Deionized water (DI water with $18.2 \text{ M } \Omega \text{ cm}^{-1}$, <5 ppb total organic carbon) was obtained from a Millipore water system.

A.1.2 Catalyst ink preparation

The Ag catalyst ink contained 200 mg of Ag in 10 mL of isopropyl alcohol and 45 μL of Nafion D-2020 dispersion. All stock solutions were sonicated for 1 hour before being passed to an automated spray-coater for deposition onto a gas-diffusion electrode.

A.1.3 Gas-diffusion electrode fabrication

After preparing the stock solution, the catalyst ink was automatically sprayed onto a $10 \text{ cm} \times 10 \text{ cm}$ H23 hydrophobic carbon paper. A total of 2 mL of the metal-containing solution was sprayed onto a pre-heated carbon paper substrate. The hotplate temperature was maintained at 150°C and was measured using a thermocouple mounted to the surface of the hotplate aluminum fixture. For details on how the thermocouples maintained the hot-plate temperature, see previous study.[21] When spraying, the nozzle moves in a serpentine pattern consisting of 50 lines 105 mm in length and evenly spaced within the 105-mm-wide carbon paper. This pattern was repeated for 10 passes with a 2 s delay between each pass. Compressed air was fed to the nozzle at 60 PSI but is restricted by an electronic air control valve (EV-P-20-2550, Clippard, USA) which was kept open at 75% of the maximum. A valve setting of 100% was measured to be 22 standard cubic feet per hour, while a value of 75% was measured to be 16 standard cubic feet per hour. The height of the nozzle above the substrate, controlled by the motorized XYZ gantry system, was kept at 5 mm. The spray flow rate, controlled by the syringe pump, was kept at 2 mL s^{-1} . After spray-coating, the samples were left to rest on the hotplate for a minimum of 60 s. The target loading of Ag nanoparticles was 1.5 mg cm^{-2} and the automated spray coating produced the gas-diffusion electrodes with a loading of $1.52 \pm 0.74 \text{ mg cm}^{-2}$. The $10 \text{ cm} \times 10 \text{ cm}$ gas-diffusion electrodes were then cut into $2 \text{ cm} \times 2 \text{ cm}$ smaller gas-diffusion electrodes, ready to be tested in ATC.

A.1.4 Automated Spray coater

The automated spray coater station is a customized 4-axis robot with an ultrasonic spray head (Supplementary Fig. 1). The ultrasonic spray coater nozzle extracts the requested volume of solution from a vial and sprays it onto a carbon paper substrate with user-defined parameters such as air flow rate, number of passes, number of raster lines, spray height, solution flow rate, and temperature. The ultrasonic spray nozzle operates at 3 W and 120 kHz. For more details about the components of the automated spray coater, please refer to previous studies.[34]

A.1.5 Automated test cell (ATC)

We upgraded the previously reported automated gas-fed CO_2 electrolyzer to perform bicarbonate electrolysis.[21] The automated test cell for high-throughput bicarbonate electrolysis testing (Supplementary Fig. 2) consists of an automated reagent (liquid and gas) delivery system, a catholyte temperature controller (Supplementary Fig. 3), and a power supply unit. The ATC was designed to have all the elements of a laboratory-scale zero-gap manual bicarbonate electrolyzer (Supplementary Fig. 4). The gas-diffusion electrode cathode was manually placed in the electrochemical chamber of the ATC along with a nickel foam anode and a bipolar membrane. This configuration separates the cathode and anode compartments similar to a traditional electrolyzer. Once the loading procedure

was complete, automated electrochemical testing was initiated with a single button press. A hydraulic system automatically closes and seals the cell.

The ATC was also connected to an atmospheric pressure chemical ionization triple-quadrupole mass spectrometer for quantitative measurements of gaseous products generated at the cathode (Supplementary Fig. 2).[21] A seven-point mass spectrometry calibration curve was separately created for CO and CO₂. The data acquisition, data processing, and the calculation of the Faradaic efficiency of CO production (FE_{CO}) was performed using automated Python scripts. This allows for FE_{CO} , CO₂ utilization, J_{CO} , and reactor yield values to be available immediately. For other technical details, please refer to the previous publication on the AdaCarbon platform.[21]

A.1.6 ATC reactive carbon electrolysis sequence

The ATC bicarbonate electrolysis sequence was as follows:

Priming: The anodic side of the test cell was primed with the KOH solution at a flow rate of approximately 30 mL min⁻¹. The cathodic side was primed with the KHCO₃ solution at a flow rate of approximately 90 mL min⁻¹. After priming, the peristaltic pumps were turned off during the baseline procedure.

Baseline 1: Argon carrier gas flow was turned on, purging the system of air and residual gasses from previous experiments. A flow of methane gas at 2 mL min⁻¹ was used for standardization. This step ran for a fixed amount of time that was chosen to allow the mass spectrometer signals for N₂, CO₂, CO, methane and Ar to reach a steady state.

Electrolysis: KOH and KHCO₃ electrolyte flows were resumed. The potential was applied to the cell and electrolysis began. This step ran for a fixed amount of time to allow the voltage and the mass spectrometer signals for the measured products to stabilize.

Baseline 2: The electrolyte flow and potential applied to the cell was halted. Argon and the methane standard gas flows were maintained. The length of this step was 10% of the initial baseline time to allow for mass spectrometry signals to return to baseline values. This provided an indication if the baseline of the system had changed during electrolysis.

CO calibration: A flow of CO at 2 mL min⁻¹ was maintained for 3 min to establish a two-point correction curve to quantify the amount of CO produced during electrolysis. The detector used to detect CO was not stable enough to use a previous calibration and therefore a correction is applied for this drift in the detector.

Purge: The KOH solution was pumped out of the lines and all gas flows were stopped.

A.1.7 ATC validation for reactive carbon electrolysis

We validated our ATC against a manually assembled bicarbonate electrolyzer used in our research group (Supplementary Fig. 4). The manual electrolyzer in this study consisted of a carbon composite cathode and a Ni foam anode sandwiched between cathodic and anodic serpentine flow plates. The cathode was prepared by automated spraying of the catalyst ink onto a 2.5 cm × 2.5 cm piece of H23 carbon paper to form a gas diffusion electrode. The active area of the gas diffusion electrode was the same as tested in the ATC (4 cm²). The reduction of CO₂ occurred at the cathode, forming CO and OH⁻. [29] The cathode and anode compartments were separated by a bipolar exchange membrane, which facilitates water splitting and transport of OH⁻ to the anode and H⁺ ions toward the cathode. The cathode was fed with a 3 M KHCO₃ solution at a flow rate of 100 mL min⁻¹, while the anode was fed with recirculated 1 M KOH at a flow rate of 30 mL min⁻¹.

Galvanostatic experiments were performed in triplicate at 100 mA cm⁻² for 5 minutes. The gaseous products were measured with an in-line gas chromatograph connected to the cathode compartment after electrolysis.[6, 35] The gas chromatograph (Perkin Elmer, Clarus 580) was equipped with a packed MolSieve 5 Å column, and a packed HayeSepD column. The gas chromatograph was used to detect CO and H₂ using a flame ionization detector, and a thermal conductivity detector, respectively. The Faradaic efficiencies (FE) of each product generated were calculated using mole fractions of the H₂ and CO from the gas chromatography data.

For ATC validation, three manufactured silver-based cathode gas-diffusion electrodes were tested in triplicate in both the manual electrolyzer and the ATC at 100 mA cm^{-2} . To quantify the accuracy of our ATC, we compared the absolute difference of FE_{CO} and CO_2 utilization when using the ATC and manual cell. We found that the absolute error, defined as the difference in mean values between the ATC and the manual cell, was 3.5% for FE_{CO} and 4.2% for CO_2 utilization. To quantify the precision of our ATC, we compared the standard deviation of triplicate measurements of the FE_{CO} and CO_2 utilization. We found that the standard deviation, expressed as a relative error, for FE_{CO} and CO_2 utilization were 2.08% and 1.45%, respectively. As a result, we concluded that our ATC platform yields similar results to a manual electrolyzer setup, with minimal variance between runs.

A.1.8 Electrolysis protocol

In each of our experiments, to avoid any effects of material degradation from influencing the data, we used fresh bipolar membranes, Ni foam, and silver gas-diffusion electrodes. During electrolysis, the cathode was supplied with bicarbonate catholyte, while the anode was fed with recirculated KOH solution. Each gas-diffusion electrode sample was tested for 5 minutes in the ATC. Afterwards, the reactor yield, J_{CO} , and CO_2 utilization were determined.

A.1.9 Faradaic efficiency and CO_2 utilization

We measured the CO and H_2 selectivity of the electrolyzer at a constant current density by quantifying their respective concentrations using mass spectrometry or gas chromatography.

The Faradaic efficiency of a gaseous product k (FE_k) was determined by:

$$FE_k = \frac{n_k F \chi_k F_m}{I}, \quad (4)$$

where n_k is the number of electrons exchanged, F is Faraday's constant ($F = 96,485 \text{ C mol}^{-1}$), χ_k is the mole fraction of gaseous product k (CO or H_2) in the gaseous mixture analyzed, F_m is the molar flow rate in mol s^{-1} , and I is the total current in A.

The molar flow rate was calculated from the volume flow rate F_v using the relation:

$$F_m = \frac{p F_v}{RT}, \quad (5)$$

with p being the atmospheric pressure in Pa, R the ideal gas constant ($8.314 \text{ J mol}^{-1} \text{ K}^{-1}$), and T the room temperature in K, which is 298 K in our study.

The CO_2 utilization (U_{CO_2}) was calculated by:

$$U_{\text{CO}_2}(\%) = \frac{[\text{CO}]}{[\text{CO}_2] + [\text{CO}]} \times 100\% \quad (6)$$

where $[\text{CO}]$ and $[\text{CO}_2]$ represent the concentration of product CO produced during electrolysis and CO_2 in the catholyte headspace outlet as measured by in-line gas chromatographic analysis.

A.1.10 Optimization campaign initialization

A set of 12 initial random experiments were chosen to initialize the Bayesian optimization. This amount was chosen based on an arbitrary $2n$ rule, where n is the number of dimensions (in our case $n = 6$). A random value between the set range of each parameter (Supplementary Table 1) was chosen for each experimental condition from a uniform random distribution (Python `numpy.random.random`). This value was then scaled to the range of the variable.

A.1.11 Reactor yield optimization campaign

To choose the next experiment after the initialization, we employed a Bayesian optimization algorithm that maximized the reactor yield. The Bayesian optimization was performed using the Botorch Python package[36] and happened in two steps. First, a surrogate Gaussian process regression

model was built over all existing data to predict the reactor yield of potential formulations using electrolysis measurements from a limited dataset. The data covariates were normalized to the unit cube and outcomes were standardized (zero mean, unit variance). The surrogate model predicts the experimental output (reactor yield) based on inputs such as bicarbonate solution temperature, concentration, flow rate, current density, KOH concentration, and flow rate. We implemented the Gaussian process model using the GPyTorch Python package.[37]

A fixed-noise Gaussian process model was used for the surrogate model.[27] We chose 3 mA cm^{-2} as the noise estimate for the Gaussian process model, which comes from the standard deviation from three independent runs, performed under the same experimental condition in the initialization campaign (see Supplementary Note 6). All other hyperparameters for the Gaussian process model were defaulted. The model prediction performance using leave-one-out cross validation is discussed in the Supplementary Note 7 and Supplementary Fig. 13.

Next, the best acquisition function was selected based on the simulation of oracle surface results (Supplementary Note 2), and a new set of experimental conditions were acquired. We selected a q-noisy expected improvement (qNEI) function[28] (Supplementary Figs. 6–7). All the sampling strategies were configured to maximize the reactor yield while also effectively searching the parameter space. The qNEI function helps manage the balance between exploration (probing new regions of the parameter space) and exploitation (refining existing knowledge where improvements are deemed likely) during Bayesian optimization. This is achieved by modulating the selection of sampling points based on both the predicted improvement and the associated uncertainty, thus mitigating the risk of redundantly sampling the same points. The inclusion of the “Noisy” term accounts for the inherent variability and uncertainty in the electrolysis experimental measurements, a feature not typically addressed by the standard expected improvement and upper confidence bound functions (see Supplementary Note 2 and Supplementary Figs. 6–7). This feature allowed the algorithm to better navigate the trade-off between exploring new, uncertain regions and exploiting areas known to yield high outcomes, ensuring robustness against experimental noise and reducing the likelihood of repetitive sampling in less promising areas.

In order to ensure a comprehensive exploration of the parameter space and test a few human-driven hypotheses, the results from some researcher-defined (RD) experiments were introduced to the optimizer (Supplementary Note 3 and Supplementary Table 2). First, the space-filling point selected an experimental condition in the parameter space that maximized the distance to the closest other existing experiment.[34] This strategy enhanced exploration by testing a few hypotheses and reducing the risk of entrapment in local optima. This ensured a comprehensive search of the parameter space.

Once a set of experimental conditions was identified using the qNEI algorithm, the optimizer provided a set of parameters for new experimental conditions to be tested in ATC. The closed loop experiments were repeated until no improvement in the reactor yield objective function was observed over a minimum of 10 consecutive steps, as shown in Fig. 2A. This stabilization was confirmed by reviewing the contour plots of the predicted outcomes and observing how these predictions changed as the optimization advanced (Supplementary Note 8 and Supplementary Fig. 8).

A.1.12 J_{CO} optimization campaign

For maximizing J_{CO} , we fed all the data obtained in the reactor yield optimization campaign as an initialization into our optimizer. The objective function selected this time was J_{CO} , as opposed to reactor yield. The optimization process was similar to the one performed in the reactor yield campaign, except that we did not include any space-filling or RD points during this experimental campaign. This block of optimization relied solely on the extensive initial dataset to guide the optimization trajectory. The absence of additional gains in J_{CO} after over 10 consecutive samples (Fig. 2B) substantiated the decision to cease further experimentation.

By starting with a denser initial dataset, the Gaussian process model used in Bayesian optimization can generate more accurate and reliable estimations of the predicted outcomes across the parameter space. A more comprehensive initial dataset reduces the uncertainty in these estimations because the Gaussian process model has access to a wider range of data points that inform its predictions. This improved estimation capability directly impacts how the acquisition function—such as the qNEI—is calibrated. With more precise estimations of predicted outcomes, the acquisition function can more effectively discern between areas of the parameter space that are likely to yield improvements and those that are still uncertain and thus require further exploration. This refined calibration helps in

making more informed decisions about where to sample next, enhancing the overall efficiency and effectiveness of the optimization process.

For example, as shown in Fig. 2B, initially the model favored exploitation, leveraging high confidence in the parameter space derived from the comprehensive initial data to hone in on areas predicted to yield high performance. After this phase of intense exploitation, the model shifted to explore a new region, in response to less certain areas in the existing parameter space. Following this exploratory point, the optimizer returned to exploitation, discovering another peak performance area (J_{CO} of 140 mA cm^{-2}). After the initial sequence of four points focused on exploitation and one on exploration, the model adopted a more balanced approach, oscillating between exploring new possibilities and exploiting known high-yield regions, thereby optimizing the search for global optima.

A.2 Supplementary Figures and Tables

This section contains supplementary figures 1-13 and supplementary tables 1-4.




Figure S1: **Automated catalyst spray-coating station as a part of AdaCarbon platform.** The automated spray-coater station is a purpose-built 4-axis robot with an ultrasonic spray head. The ultrasonic spray coater nozzle extracted the required amount of solution through a needle from the vial and commenced spraying onto the 10 cm $\times$ 10 cm Freudenberg H23 carbon paper with user-defined parameters such as air flow rate, number of passes, number of lines, spray height, spray flow rate and temperature of the hot plate. (the glass slide is only shown as an example). The ultrasonic spray nozzle operated at 3 W and 120 kHz.




Figure S2: Automated test cell station for performing liquid-fed bicarbonate electrolysis as a part of AdaCarbon platform. (A) Photograph of an Automated Test Cell (ATC) platform developed in this study. Instead of relying on a manual assembly and disassembly process, the cell is sealed automatically by a hydraulic system once the researcher loads a GDE into it. During electrolysis validation, a 1 M KOH solution is circulated through the anode, facilitated by a computer-controlled peristaltic pump. A 3 M KHCO_3 , is pumped to the cathode from a catholyte reservoir, connected with the automated temperature controller. The resulting product stream is then subjected to a gas dryer and analyzed by a mass spectrometer. (B) Models of the automated test cell (ATC) for performing bicarbonate electrolysis and its comparison with the manually assembled electrolyzer. The ATC is structured similarly to a manually assembled electrolyzer using the same stack of flow plates, cathode, anode, and ion-exchange membrane. However, the ATC has no assembly fasteners; instead, it uses a hydraulic piston to close the cell automatically. The raw data acquisition and processing, including experimental conditions, is completely automated for the ATC system.




Figure S3: **Automated catholyte temperature controller reservoir installed in ATC.** (A) The CAD model of the bicarbonate catholyte reservoir (with pre-existing fittings) was employed in this study. It uses a thermocouple to sense the accurate measurements of the temperature and has an immersion heating system which requires minimal hardware and has inherent insulation benefits. It contains a valve in the electrolyte reservoir which could be automatically relieved. (B) The automated temperature controller was used in this study. The leads of the instrument are directly connected to the thermocouple in the bicarbonate reservoir. A user can set the temperature manually. The temperature of the electrolyte reservoir can be controlled within $\pm 3\text{ }^{\circ}\text{C}$ for the duration of an experiment. A temperature range between 25–80 $^{\circ}\text{C}$ can be achieved. If a temperature $>100\text{ }^{\circ}\text{C}$ is detected, the heating apparatus automatically shuts off.




Figure S4: Manual Bicarbonate electrolysis setup. A manual cell consists of a gas-diffusion cathode and Ni foam anode separated by a bipolar exchange membrane. A power supply is connected to the cathode and anode terminals to perform electrolysis in galvanostatic mode. A peristaltic pump transports 1 M KOH at a desired flow rate through the anode flow plate to a nickel foam anode (unless stated). The flow is recycled to the 1 M KOH reservoir and O₂ is vented. The cathode flow plate is fed with a 3 M KHCO₃ solution (unless stated). The gas flow is controlled by a mass flow controller (FC). The mixture of gaseous products (CO, CO₂, and H₂) is passed through the condenser which then is allowed to pass through a gas chromatogram for FE_{CO} , J_{CO} , U_{CO_2} and reactor yield quantification.




Figure S5: **Number of points vs. number of variables tested.** For this specific search space with 6 dimensions, there are about 15,600 possible combinations. To elucidate the computation of the 15,600 combinations mentioned, consider the method of calculating the number of potential settings for each parameter. For each of the six dimensions, we assume that the researcher selects 5 discrete points within the total range of the parameter, which gives 5 possible settings for each parameter. The product of these individual counts across all parameters yields the total number of possible combinations. Multiplying these numbers together ($5 \times 5 \times 5 \times 5 \times 5 \times 5$) provides the total number of combinations, which is 15,625. This multiplication reflects the exponential growth of the search space with each added dimension and the corresponding discrete settings, showcasing the massive scale of our exploration space.




Figure S6: **Q-noisy expected improvement function benchmarking against alternative sampling strategies.** The performance of the q-noisy expected improvement algorithm is compared to alternative single-objective sampling strategies (namely the qEI and qUCB algorithms, random sampling) in simulated optimization campaigns for (A) reactor yield and (b) J_{CO} . The median (solid line) and interquartile range (shaded bands) from replicate simulations are shown. 50 replicate simulations were performed for all methods except random, for which random where 100,000 replicates were performed.




Figure S7: **Quantitative assessment of the improvement provided by the q-noisy expected improvement algorithm in the simulated optimization campaign.** The enhancement factor achieved by the simulated q-noisy expected improvement algorithm relative to random searches is depicted, based on the (A) reactor yield and (B) J_{CO} optimization experimental data. The x-axis represents the samples, which correspond to the experiment numbers used throughout the manuscript. For the simulations, the median results are shown with a solid line, the interquartile range is indicated by shaded bands, and the geometric mean is represented by a dashed line.




Figure S8: Contour plots of the posterior mean at three stages: after initialization, after the first optimization campaign, and after the final optimization campaign. Contour plots are presented for two pairs of parameters: (A) KOH concentration vs. bicarbonate flow rate, and (B) bicarbonate flow rate vs. current density. In the plots, black points represent the initialization samples, orange points represent all samples collected after initialization, including optimization samples, and the red star indicates the champion point for the reactor yield optimization campaign.




Figure S9: **Variation of reactor yield with (A) CO_2 utilization and (B) J_{CO}** “Other” represents initialization and other researcher-defined (RD) points.




Figure S10: **Variation of six parameters when optimizing for reactor yield.** Six input parameters were manipulated: four electrolyzer parameters (Current density, bicarbonate flow rate, KOH flow rate, bicarbonate temperature) and two chemical variables (KOH concentration and bicarbonate concentration). The experimental conditions requested for each experiment are shown in panels A–F. The bounds of the experimental variables are shown as the minimum and maximum y-axis values for each plot (see also Table S1). The orange line represents the input values associated with the running best reactor yield.




Figure S11: **Variation of six parameters when optimizing for J_{CO} .** Six input variables were manipulated: four electrolyzer parameters (Current density, bicarbonate flow rate, KOH flow rate, bicarbonate temperature) and two chemical variables (KOH concentration and bicarbonate concentration). The experimental parameter optimized for is shown in each panel (A–F). The bounds of the experimental variables are shown as the minimum and maximum y-axis values for each plot (see also Table S2). The orange line represents the input values associated with the running best J_{CO} .




Figure S12: **Variation of FE_{CO} versus input current density.** During the J_{CO} optimization where the initialization points came from the reactor yield optimization process, the optimizer successfully overcame the typical trend of lower J_{CO} values at higher current densities, usually due to mass transport limitations. The Bayesian optimizer helps achieve high FE_{CO} values even at elevated current densities.



A**B**

Figure S13: **Leave-One-Out Cross-Validation (LOOCV) and model residuals of the experimental response surfaces for the reactor yield.** (A) Model residuals plotted as a function of the model output: the reactor yield. The residuals from the LOOCV analysis are shown against the predicted reactor yield values. The coefficient of determination (R^2) for the model is 0.86, indicating a high level of predictive accuracy. (B) LOOCV analysis of the experimental response surfaces for the reactor yield. The Gaussian process model was trained on all the combined data from the reactor yield optimization campaign.

Table S1: **A list of six electrolyzer parameters manipulated in this work.** The parameter range and importance are also mentioned.

Parameters	Range	What does this parameter influence?
Bicarbonate temperature (°C)	25–80	solubility of <i>i</i> -CO ₂
Bicarbonate concentration (M)	0.5–3.05	source of <i>i</i> -CO ₂
Bicarbonate flow rate (mL min ⁻¹)	30–200	residence time of <i>i</i> -CO ₂
Current density (mA cm ⁻²)	50–350	CO selectivity
KOH concentration (M)	0.05–3	cathodic chemistry
KOH flow rate (mL min ⁻¹)	5–30	cathodic chemistry

Table S2: **An overview of researcher-defined (RD) points.**¹ The first condition is for testing a human-driven hypothesis with points selected from manual cell experiments. The second and third points relate to testing the effect of low and high bicarbonate concentration on reactor yield, keeping all other variables at standard conditions. The fourth and fifth points relate to varying bicarbonate concentrations for champion reactor yield conditions. The champion reactor yield was obtained at the bicarbonate concentration of 0.52 M. The last five points came from the space filling algorithm.

No.	Bicarbonate temperature (°C)	Bicarbonate concentration (M)	Bicarbonate flow rate (mL min ⁻¹)	Current density (mA cm ⁻²)	[KOH] (M)	KOH flow rate (mL min ⁻¹)	Reactor yield (mA cm ⁻²)
1	40	1.5	150	200	1.5	30	23.67
2	25	3.0	100	100	1.0	30	11.01
3	25	0.5	100	100	1.0	30	12.99
4	55	3.0	106	249	1.7	24	13.83
5	55	1.5	106	249	1.7	24	17.77
6	61	2.3	187	332	0.1	28	24.32
7	68	2.5	88	184	2.5	7	15.19
8	31	2.7	52	300	1.7	20	13.38
9	34	1.7	54	110	1.4	28	9.11
10	74	1.0	54	213	0.6	29	16.55

Table S3: **The list of parameters corresponding to the highest reactor yield and J_{CO} values obtained from their respective optimization experiments.**

Parameters	Maximized reactor yield	Maximized J_{CO}
Bicarbonate concentration	0.52 M	2.37 M
Bicarbonate flow rate	144 mL min ⁻¹	192 mL min ⁻¹
Bicarbonate temperature	54 °C	63.8 °C
KOH concentration	2.4 M	0.19 M
KOH flow rate	24 mL min ⁻¹	28.1 mL min ⁻¹
Current density	249 mA cm ⁻²	329 mA cm ⁻²

Table S4: **Parameters for demonstration of the practical scenario.** The total concentration of the K^+ ions was kept constant at 3 M. All other parameters were constant as obtained for the champion condition for the reactor yield. Bicarbonate flow rate = 144 mL min^{-1} , bicarbonate temperature = $54 \text{ }^\circ\text{C}$, KOH concentration = 2.4 M, KOH flow rate = 24 mL min^{-1} , and current density = 249 mA cm^{-2} .

$[\text{HCO}_3^-]$ (M)	$[\text{CO}_3^{2-}]$ (M)	U_{CO_2} (%)	J_{CO} (mA cm^{-2})	Reactor yield (mA cm^{-2})
0.0	1.50	100.0	17.3	17.3
0.5	1.25	98.0	31.6	30.9
1.0	1.00	94.0	42.7	40.2
1.2	0.90	89.0	43.9	39.0
1.5	0.75	81.1	54.2	43.9
1.8	0.60	75.3	55.5	41.8
3.0	0.00	27.1	87.4	23.7

A.3 Supplementary Note 1: Bayesian optimization and acquisition function

Bayesian optimization is a powerful method for optimizing black-box functions that are expensive to evaluate and may include observation noise. It employs a probabilistic surrogate model, such as a Gaussian Process (GP), to model the objective function $f(x)$ and quantify uncertainty in predictions. This surrogate model allows the algorithm to make informed decisions about where to sample next by balancing exploring uncertain areas and exploiting promising regions.

The selection of the next evaluation point is guided by an acquisition function, which measures the expected utility of sampling at a candidate point. In noisy settings, the q-noisy expected improvement (qNEI) is an appropriate choice as it extends the standard Expected Improvement (EI) criterion to account for observation noise. In this study, we use a batch of one experiment ($q = 1$).

The standard Expected Improvement (EI) at a candidate point x is defined as:

$$EI(x) = \mathbb{E} [\max (0, f_{\text{best}} - f(x))],$$

where f_{best} is the best-observed function value so far. However, when observations are noisy, we observe

$$y = f(x) + \epsilon,$$

with ϵ representing noise, and the true f_{best} is uncertain.

The Noisy Expected Improvement (NEI) addresses this by integrating over the posterior distributions of both the true function values at the observed points and the candidate point. The NEI is formulated as:

$$NEI(x) = \mathbb{E}_{f(x), f_D} [\max (0, f_{\min} - f(x))],$$

where f_D are the true function values at the observed data points D , and

$$f_{\min} = \min(f_D).$$

Under the Gaussian Process (GP) model, the joint distribution of $f(x)$ and f_D is multivariate normal. The NEI can be estimated using Monte Carlo integration:

1. Sample true function values $f_D^{(i)}$ from their posterior distribution given the noisy observations.
2. Compute the minimal true value $f_{\min}^{(i)} = \min (f_D^{(i)})$ for each sample.
3. Sample the function value at x , $f^{(i)}(x)$, from the conditional posterior distribution given $f_D^{(i)}$.
4. Calculate the improvement

$$I^{(i)} = \max (0, f_{\min}^{(i)} - f^{(i)}(x)).$$

5. Estimate NEI as

$$NEI(x) \approx \frac{1}{N} \sum_{i=1}^N I^{(i)},$$

where N is the number of Monte Carlo samples.

Alternatively, NEI can be approximated analytically under certain assumptions. When the noise is small and the GP posterior at x remains approximately Gaussian, NEI simplifies to:

$$NEI(x) \approx (f_{\text{best}} - \mu(x)) \Phi\left(\frac{f_{\text{best}} - \mu(x)}{\sigma(x)}\right) + \sigma(x) \phi\left(\frac{f_{\text{best}} - \mu(x)}{\sigma(x)}\right),$$

where

- $\mu(x)$ is the posterior mean of the GP at point x ,
- $\sigma(x)$ is the posterior standard deviation at x ,

- $\Phi(\cdot)$ and $\phi(\cdot)$ are the cumulative distribution function and probability density function of the standard normal distribution, respectively.

In this approximation, f_{best} is often taken as the minimal posterior mean among observed points.

The NEI effectively accounts for the uncertainty due to noise by considering the distribution over the true function values rather than just the observed noisy outputs. This leads to more robust optimization in noisy settings, as it prevents the algorithm from being overly optimistic about noisy observations.

A.4 Supplementary Note 2: Simulation to model the oracle surface and quantification of the algorithm performance

We employed computer simulations to evaluate the efficacy of the qNEI algorithm in the optimization of reactor yield and J_{CO} experiments. An experimental response surface model, constructed from collected data, served as the basis for these simulations. We conducted simulated optimizations using various algorithms—random search, qUCB, qEI, and qNEI—sampling the model to compare their performances (Fig. S6). The effectiveness of the qNEI algorithm, particularly against random sampling, was assessed using the enhancement factor (EF) metric (shown in Fig. S7). Detailed descriptions of these simulation procedures are further elaborated below.

A.4.1 Models of the experimental response surface

: To accurately model the experimental conditions, Gaussian process (GP) regression was employed to develop two distinct response surface models (oracles), each tailored to one of the objectives: reactor yield and J_{CO} . These models predict the experimental outputs based on inputs such as bicarbonate temperature, concentration, current density, flow rate, KOH concentration, and flow rate. We implemented the GP response surfaces using the GPyTorch Python package.² For each model, normalization of the input data, as well as standardization (zero mean, unit variance) of the respective objective (reactor yield or J_{CO}), was performed before training each model (response surface).

A.4.2 Sampling strategies

:

For evaluating the efficacy of various sampling strategies, we conducted simulations using different sampling strategies (random search, qUCB, qEI, and qNEI algorithms) on the reactor yield and J_{CO} experimental models. Each simulation utilized identical initialization points across the reactor yield and J_{CO} setups consistent with the corresponding reactor yield and J_{CO} experiments to ensure comparability. Random search involved generating samples from a uniform distribution throughout the normalized input space. The qUCB algorithm was configured with a beta value of 1 (beta values can control the exploration of space), chosen to balance exploration and exploitation effectively within this context, ensuring moderate confidence in the uncertainty bounds of the predictions. Both qEI and qNEI algorithms were configured using BoTorch's default settings. The complete benchmarking results are illustrated in Fig. S6 and Fig. S7, providing a visual comparison of the performances of these algorithms.

A.4.3 Simulated optimization campaigns

:

Each simulated optimization campaign was performed for 50 replicates and 100 experimental iterations, except for random which was performed for one hundred thousand iterations for use in the enhancement factor calculations (shown in Fig. S6 and Fig. S7). If an optimization algorithm produced an error during optimization, then that replicate was removed and repeated.

A.4.4 Enhancement Factor Calculation

: The enhancement factor for a given sampling technique, relative to another, is defined as the ratio of their performance values when the same number of samples are used. For instance, if sampling method A achieves a performance value of 10 with 50 samples, and method B achieves a value of 5 with the same 50 samples, then the enhancement factor of technique A at 50 samples is 2. The

calculation follows the equation below:

$$EF_{A:B}(n) = \frac{P_A(n)}{P_B(n)},$$

where $EF_{A:B}(n)$ is the enhancement factor of sampling technique A to B after n samples. When $P_B(n) = 0$, then $EF_{A:B}(n)$ is not calculable.

To compare the enhancement factor of the qNEI algorithm relative to random search across 50 replicate simulations, the median, geometric mean, and interquartile range of the enhancement factor were computed and illustrated in Fig. S7.

A.5 Supplementary Note 3: Reactor yield optimization strategy and interventions by researcher

The reactor yield optimization campaign commenced with an initialization phase comprising 12 experiments using randomized parameters. Subsequently, 10 experiments utilizing Bayesian optimization (BO) followed, increasing reactor yield from 31 mA cm^{-2} to 33.5 mA cm^{-2} by the 13th experiment. However, the optimizer concentrated on exploiting regions around previously high reactor yield values, potentially constraining broader parameter space exploration.

To address this, two targeted interventions were introduced by the researcher, which we call *researcher-defined (RD) points*:

1. **Space-filling points for enhanced exploration:**

After the 22nd experiment, 5 space-filling (SF) points³ were added to maximize diversity in the input space, steering the optimization away from previously explored regions. This intervention broadened the search and facilitated a restart of the optimization process. Consequently, a new maximum reactor yield of 38 mA cm^{-2} (the champion condition) was achieved after 9 additional experiments, demonstrating the effectiveness of combining space-filling exploration with targeted optimization.

2. **Human-in-the-loop (HIL) interventions:**

Following the discovery of the champion condition, the optimizer trended toward lower bicarbonate concentrations. To counter this, 5 human-in-the-loop points were added, focusing on two objectives:

- (a) Encourage exploration by varying bicarbonate concentrations (0.5 M, 1.5 M, 3 M) under both champion and standard conditions.
- (b) Test researcher-driven hypotheses to balance CO_2 utilization and J_{CO} (Table S2).

Although none of the HIL points outperformed the optimizer, they validated process stabilization and tested some hypotheses. This was confirmed by examining contour plots of predicted outcomes and their evolution throughout optimization (see Fig. S13).

A.6 Supplementary Note 4: SHAP Analysis

To explain the underlying Gaussian process models that drive Bayesian optimization for reactor yield and J_{CO} , we used the SHAP (Shapley Additive exPlanations) framework,⁴ which is based on a coalitional game-theoretic approach. We first fit the entire dataset using a random forest (RF) regressor model, which constructs a multitude of decision trees.⁵

SHAP was then used to identify the important features that best explain the output of the RF model and gain additional physical insights. The major goal of SHAP is to explain the prediction (or instance) of the black-box decision of an ML model by a stepwise decomposition of the model predictions and compute the contribution of each feature (or group of features) to the prediction. The SHAP values for a specific data point (which is an experiment number in our case) are generated by comparing the model outcome difference of including or not including a certain parameter or feature.

To compute SHAP values, a baseline or reference value is determined for each feature, representing its expected or average value in the dataset. The SHAP value of a feature indicates how much its specific value in a prediction deviates from the baseline. In essence, it quantifies the feature's contribution to shifting the model's prediction away from the average or expected output.

The features of an instance are treated as players in a coalition, with the coalition representing the set of features involved in a particular prediction. Using SHAP values to interpret an ML model allows us to understand how each feature (or player) contributes to the prediction. The phrase “players in a coalition” refers to the individual features or variables considered when calculating SHAP values for a specific prediction or outcome.

SHAP can provide both local (for a single data point, a single surface in our case) and global (takes into account the entire dataset, and all experiments) explanations for the predictions of an ML model.⁴ A so-called beeswarm SHAP summary plot,⁴ which provides a global explanation of the RF model predictions (Figs. 3A and 3B in the main text). For each feature (or parameter), the SHAP values for the entire dataset are plotted as points along the horizontal axis. In positions with a high density of SHAP values, the points are vertically stacked. The global importance of each feature in prediction is ranked according to its mean absolute SHAP value, which is averaged over the complete dataset.

The horizontal axis (Figs. 3A and 3B in the main text) shows the SHAP values, where a value of zero represents no contribution of that particular feature to the prediction; contributions increase in either direction as the SHAP value moves away from zero. Whether a SHAP value is positive or negative depends on how the feature’s value in the prediction compares to its expected or average value as discussed above. The width of the distribution of points along the horizontal axis gives a qualitative indication of how strongly that particular feature impacts the model prediction. The vertical color bar depicts the absolute feature value (in our case - electrolysis parameters) from low (blue) to high (red). We emphasize that the color indicates the “raw” value of the feature in the dataset.

A.7 Supplementary Note 5: Demonstration of practical experimental scenario

As a practical scenario, the bicarbonate electrolyzer needs to be connected with the CO₂ capture plant. The CO₂ capture plant has a mixture of carbonate (CO₃²⁻) and bicarbonate ions (HCO₃⁻) in varying ratios as an eluent. To demonstrate this experimentally, we selected the champion parameters that achieved the highest reactor yield (see Table S3) and varied the composition of the capture solution. We only varied the bicarbonate concentration in the carbonate/bicarbonate mixture from 0 to 3 M and supplemented the rest of the K⁺ ions coming from the carbonate solution while maintaining a total K⁺ concentration of 3 M (see Table S4). We show that for a bicarbonate electrolysis technology to be a viable solution (measured by reactor yield), we only require 1–1.8 M bicarbonate concentration in the eluent as opposed to 3 M, which has been the focus of most of the previous studies.^{1,6,7}

A.8 Supplementary Note 6: Source of variability in the experiments

The standard deviation for reactor yield and J_{CO} values for three independent experiments were found to be 3 mA cm⁻². There are several possible sources of variations in these metrics. These sources include variations in the loading of Ag on GDEs during automated spraying (caused due to variable thermal contact of GDEs on the hot plate); variation in manual dilution of the catholyte and anolyte solutions, or variability in the ion exchange membranes from batch to batch.

A.9 Supplementary Note 7: Leave-one-out cross-validation (LOOCV) analysis of the experimental response surface

In the LOOCV analysis, the model’s accuracy is evaluated by systematically leaving out one data point at a time from the dataset. The model is then trained on the remaining data to predict the outcome for the excluded point. This process is repeated for each data point, providing an overall measure of the model’s performance. The residuals, which represent the differences between the predicted and actual objective values (in this case, the reactor yield), are used to assess how accurately the model captures the experimental results. These residuals are plotted against the predicted reactor yield to highlight any discrepancies. Additionally, the predicted reactor yields are compared with the measured values, and the coefficient of determination (R^2) is calculated (Fig. S12).

A.10 Supplementary Note 8: Progress of the optimizer using the contour plots and quantifying posterior mean

To observe how the optimizer progresses over the reactor yield optimization experiment, we plotted a contour plot for each combination of the 6 input parameters to give us 15 pair plots (results for two

pair plots are shown in Fig. S13 as examples). The contour plots illustrate the dynamic evolution of the posterior mean (expected outcome) throughout the optimization experiments for various parameter configurations relevant to electrolysis.

Initially, the distribution of initialization points (black) spans a broad expanse of the parameter space, setting a baseline for the response surface. As optimization proceeds, the insertion of points derived from optimization algorithms (orange) progressively refines our understanding, focusing on areas of the parameter space suggestive of higher yields, as indicated by shifts in the contour lines and intensification of colors towards warmer hues in specific plots (Fig. S13).

The decision to conclude the optimization campaign was based on observed diminishing marginal returns in the adjustments to the posterior mean, signaling convergence toward an optimal set of parameters. The strategic inclusion of SF and HIL points, systematically integrated after every ten optimization experiments, facilitated enhanced exploration and minimized the risk of local optima entrapment, ensuring a thorough exploration of the parameter space. This approach, coupled with the evident stabilization in the evolution of the posterior mean plots after each block — comprising initialization, optimization, space-filling, and HIL stages — provided robust empirical justification for terminating further experimentation, as continued efforts were projected to yield minimal additional gains.

A.11 Appendix Acknowledgments

We acknowledge the efforts of the contributors to the many open-source software packages used in this work:

- Python (<https://www.python.org/>)
- NumPy (<https://numpy.org/>)
- pandas (<https://pandas.pydata.org/>)
- Matplotlib (<https://matplotlib.org/>)
- SciPy (<https://www.scipy.org/>)
- scikit-learn (<https://scikit-learn.org/>)
- Plotly (<https://plotly.com/>)
- BoTorch (<https://botorch.org/>)
- BoTorch License (<https://github.com/pytorch/botorch/blob/main/LICENSE>)
- PyTorch (<https://pytorch.org/>)
- Git (<https://git-scm.com/>)

A.12 Appendix References

1. Li, T.; Lees, E. W.; Goldman, M.; Salvatore, D. A.; Weekes, D. M.; Berlinguette, C. P. Electrolytic Conversion of Bicarbonate into CO in a Flow Cell. *Joule* **2019**, 3 (6), 1487–1497.
2. Gardner, J. R.; Pleiss, G.; Bindel, D.; Weinberger, K. Q.; Wilson, A. G. GPyTorch: Blackbox Matrix-Matrix Gaussian Process Inference with GPU Acceleration. *arXiv [cs.LG]*, 2018.
3. Rupnow, C. C.; MacLeod, B. P.; Mokhtari, M.; Ocean, K.; Dettelbach, K. E.; Lin, D.; Parlange, F. G. L.; Chiu, H. N.; Rooney, M. B.; Waizenegger, C. E. B.; de Hoog, E. I.; Soni, A.; Berlinguette, C. P. A Self-Driving Laboratory Optimizes a Scalable Process for Making Functional Coatings. *CR-PHYS-SC* **2023**, 4 (5).
4. Welcome to the SHAP documentation — SHAP latest documentation. <https://shap.readthedocs.io/en/latest/> (accessed 2024-07-29).
5. Breiman, L. Random Forests. *Mach. Learn.* **2001**, 45 (1), 5–32.
6. Li, T.; Lees, E. W.; Zhang, Z.; Berlinguette, C. P. Conversion of Bicarbonate to Formate in an Electrochemical Flow Reactor. *ACS Energy Lett.* **2020**, 5 (8), 2624–2630.
7. Pimlott, D. J. D.; Jewlal, A.; Kim, Y.; Berlinguette, C. P. Oxygen-Resistant CO₂ Reduction Enabled by Electrolysis of Liquid Feedstocks. *J. Am. Chem. Soc.* **2023**, 145 (48), 25933–25937.