Accelerated High-Entropy Alloys Discovery for Electrocatalysis via Robotic-Aided Active Learning

Zhichu Ren ZC_REN@MIT.EDU
Department of Materials Science and Engineering
Massachusetts Institute of Technology
Cambridge, MA 02139-4307, USA

Zhen Zhang ZHANG21@MIT.EDU
Department of Materials Science and Engineering
Massachusetts Institute of Technology
Cambridge, MA 02139-4307, USA

Yunsheng Tian YUNSHENG@CSAIL.MIT.EDU
Department of Electrical Engineering and Computer Science
Massachusetts Institute of Technology
Cambridge, MA 02139-4307, USA

Ju Li LIJU@MIT.EDU
Department of Materials Science and Engineering
Department of Nuclear Science and Engineering
Massachusetts Institute of Technology
Cambridge, MA 02139-4307, USA

Abstract

This work explores the accelerated discovery of High-Entropy Alloys electrocatalysts using a novel carbothermal shock fabrication method, underpinned by an active learning approach. A high-throughput robotic platform, integrating a BoTorch-based active learning module with an Opentrons liquid handling robot and a 7-axis robotic arm, expedites the iterative experimental cycles. The recent integration of large language models leverages ChatGPT's API, facilitating voice-driven interactions between researchers and the automation setup, further enhancing the autonomous workflow under experimental materials science scenarios. Initial optimization efforts for green hydrogen production catalyst yield promising results, showcasing the efficacy of the active learning framework in navigating the complex materials design space of HEAs. This study also emphasizes the crucial need for consistency and reproducibility in real-world experiments to fully harness the potential of active learning in materials science explorations.

Keywords: Active Learning, Robotic Platform, High-throughput Experiment, Electrocatalysis

1. Introduction

Conventional electrocatalysts are somewhat restricted regarding compositional diversity, which may result in insufficient active sites for multi-step reactions, limited tunability for bonding with different intermediates, and a common dependency on noble metals (Peng et al., 2021). In contrast, High-Entropy Alloys (HEAs), which comprise more than five

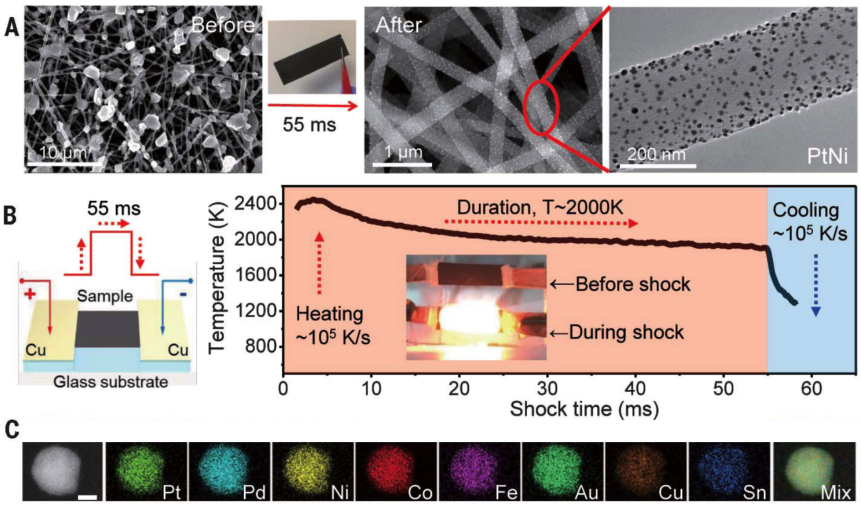

Figure 1: CTS-enabled creation of HEA nano-particles on carbon substrate. (A) Microscopic images of micro-scaled precursor salt grains on the carbon nanofiber pre-thermal shock, and the resultant, well-distributed (PtNi) nanoparticles post-CTS. (B) Sample formulation and the time-wise temperature transition during the 55-ms thermal shock. (C) Elemental mappings of an HEA nano-particle consisting of eight varied elements (Pt, Pd, Ni, Co, Fe, Au, Cu, and Sn). Scale bar, 10 nm. Reprinted with permission from AAAS.

metallic elements, have demonstrated notable improvements in both activity and durability across various electrocatalytic reactions (Yao et al., 2022). The vast domain of chemical design holds immense promise for HEAs, yet also poses a significant challenge for researchers who may struggle to intuitively navigate these high-dimensional spaces.

A recent advancement in the fabrication of HEAs catalysts has enabled a rapid iteration of recipe exploration (Yao et al., 2018). In this process, researchers merely mix the metal precursor solution according to a specified composition on a carbon substrate, followed by an application of a direct current, as shown in Figure 1. The resulting joule heat swiftly elevates the temperature to 2000K within mere tens of milliseconds, before quickly cooling down to room temperature. This intense thermal procedure is termed carbothermal shock (CTS), and the resultant product emerges well mixed at the atomic level, embodying HEAs.

Employing a typical workflow of Gaussian Process and Bayesian Optimization can guide researchers to effectively explore and rapidly converge on the global optimum within the materials design space (MDS) (MacLeod et al., 2020; Burger et al., 2020; Sun et al., 2019). The inclusion of a robot-assisted workflow further bolsters the consistency and reproducibility of experiments, establishing a solid base for the active-learning approach to function effectively. When tailored properly, this advanced research paradigm can conserve time and resources, redirecting researchers' attention from monotonous lab tasks to more intellectual analysis and design (Stach et al., 2021; Stein and Gregoire, 2019).

2. Approach

A high-throughput robotic platform consisting of four main modules, shown in Figure 2, has been developed. First, the BoTorch-based (Balandat et al., 2019) active learning module will propose the recipes for the next batch of experiments, based on the data we collected so far. These recommended recipes are then forwarded to Opentrons, a liquid handling robot, which mixes pre-prepared metal precursors, dispensing a specific volume of the mixture onto a carbon substrate. Subsequently, the CTS is carried out on the carbon substrate, which is then laser-cut into standard-sized pieces. Once the standard-sized catalyst samples are placed on the electrode sample holder, a 7-axis robotic arm manages the loading and unloading of samples into the three-electrode testing cell, while a Python script emulates human mouse and keyboard input to operate the testing software. Post-testing, the raw data is automatically analyzed to derive a metric value, which is then uploaded to a cloud SQL server for integration in the subsequent active learning cycle, thus completing the loop. This entire process is showcased in our YouTube demo (Ren, 2022).

Very recently, we implemented a huge upgrade to our high-throughput robotic platform by developing an AI research copilot on it (Ren, 2023a). Previously, autonomous labs relied heavily on scripting languages like Python, restricting their use among materials experimentalists without coding background. The advent of OpenAI's ChatGPT API's function calling feature now facilitates the integration and execution of Python subroutines in experimental workflows via voice commands. Our novel Copilot for Real-world Experimental Scientist (CRESt) system, demonstrated on YouTube, harnesses large language models (LLMs), enabling all research group members to utilize the robotic platform for their projects through simple voice interactions with CRESt.

The CRESt "operating system" comprises four main components: a user interface, ChatGPT back-end, active learning, and end-effectors. The user interface, built on chatgpt-voice (thanhsonng, 2023), facilitates voice recognition and text-to-AI-generated-voice interactions. The convenient web-based framework also allows users to continue their conversations on mobile devices even after leaving the physical lab. The ChatGPT back-end, grounded on CallingGPT (RockChinQ, 2023), translates Python functions documented in Google style docstring into a JSON format recognizable by ChatGPT, enabling their invocation whenever necessary and creating a feedback loop for immediate local execution of the suggested function with the return value sent back to ChatGPT. A wrapper layer was also customized to integrate the Ax active learning module into the CRESt workflow. The end-effectors vary widely, ranging from local or online database (e.g. Materials Project (Jain et al., 2013)) retrieval module to a set of automation subroutines ready to be invoked via HTTP requests, such as liquid handling robots, laser cutters, gas valves and pumps. While subroutine automation currently heavily relies on PyAutoGUI (asweigart) to simulate human mouse and keyboard input, we foresee this step becoming redundant as more lab equipment is expected to offer dedicated AI interfaces in the future.

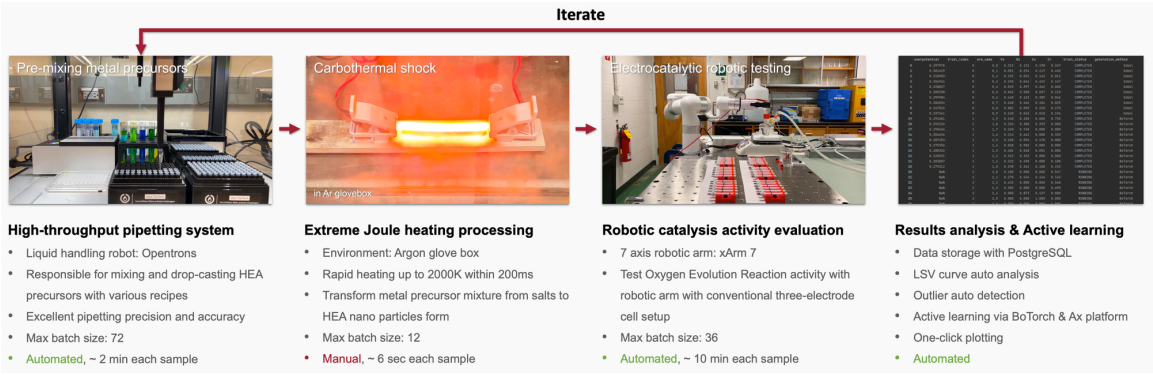

Figure 2: Pipeline of our high-throughput robotic platform with four main modules. Initially, a BoTorch module proposes next experiment batch recipes, which are sent to Opentrons for liquid handling. Following a Carbothermal Shock process and laser-cutting, standard-sized catalyst samples are managed by a 7-axis robotic arm for testing. Post-test metric values are uploaded to a cloud SQL server for future active learning cycles.

3. Results

Figure 3 shows the results from our early efforts to optimize a high-entropy alloy catalyst for the alkaline oxygen evolution reaction, a longstanding bottleneck issue in industrial green hydrogen production. The key metric for optimization is the overpotential value, with a lower value being preferable. The design space is a 5-dimensional materials domain, encompassing the composition of Fe, Co, Ni, Cr, and V, with the constraint that all compositions sum to 1. We initiated the search process with a round of SOBOL search (Sobol, 2001) and were fortunate to identify a formulation that surpassed our prior best sample (FeCoNi) discovered through manual exploration. The subsequent active learning batch employing a Gaussian Process + Expected Improvement setup further improved the record by 10mV, marking a respectable performance in this domain.

Our developed robotic platform aptly meets the requirements of electrocatalyst research. Besides the alkaline OER scenario discussed here, we also employed Gaussian Process + Upper Confidence Bound framework to hasten the investigation of catalysts for acidic oxygen evolution reaction (green hydrogen production & carbon neutrality), acidic methanol oxidation reaction (methanol fuel cell), and alkaline formate oxidation reaction (formate fuel cell). The results so far have been very promising - a particular formulation identified through active learning exhibited a performance three times better than our benchmark formulation.

4. Discussions

Upon implementing active learning in real-world materials investigation, we observed a distinct difference compared to its application in simulated or virtual environments. Typically, the outcomes of simulation are reproducible; even in scenarios with stochastic elements,

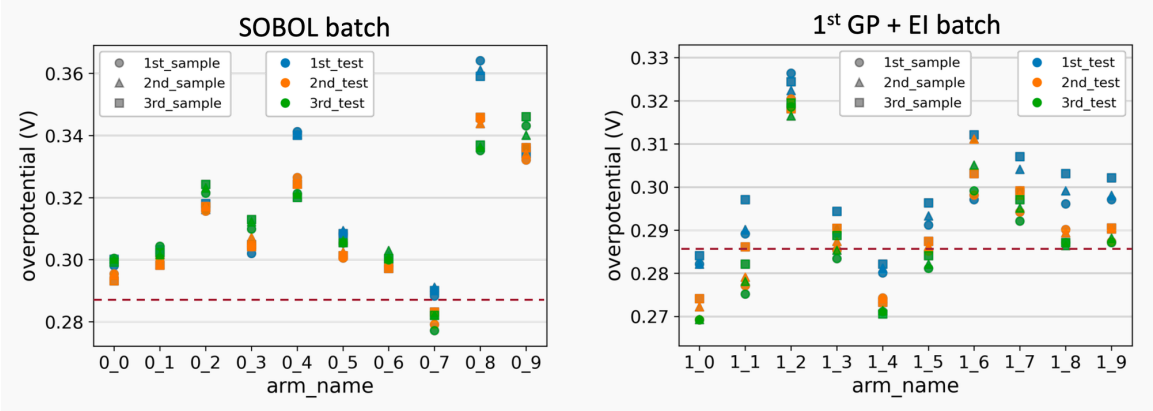

Figure 3: The initial result in green hydrogen catalyst optimization project, where lower overpotential value is preferred. The search was initiated with a SOBOL round, leading to the identification of a formulation outperforming our previous best sample (FeCoNi, red-dashed line) found through manual exploration. The following active learning batch, utilizing a Gaussian Process + Expected Improvement setup, enhanced the record by 10mV, showcasing a commendable performance in this domain.

setting a fixed random seed can mitigate result variance. This, however, isn't the case in real-world experiments. Fields like biology and materials science are infamous for their reproducibility challenges (Baker, 2016), whether temporally (inability to replicate the results from last month with the exact same input parameters) or spatially (inability to replicate the results from a different lab with the exact same input parameters).

The crux of the irreproducibility issue lies in our researchers' inability to identify all the variables involved in the experimental workflow. For instance, if an experiment workflow's outcome can be rigorously determined by 50 variables (from an omniscient perspective), but we only recognize 40 of them (either have them logged in a spreadsheet or well fixed at a constant value), while the remaining 10 variables elude our notice. In such a case, those 10 variables will be the culprit for the irreproducibility, as they may change over time or across different locations. Unfortunately, the variance induced by these hidden variables is not Gaussian, rendering them poorly handled by the noise kernel in the gaussian process.

Without meticulously inspecting the experimental workflow prior to embarking on an active learning campaign, significant resources may be wasted on an irreproducible dataset. The algorithm could mistakenly take non-gaussian noise as signal, and eventually leading to a scenario of 'garbage in, garbage out.' Hence, every experimentalist must critically consider, "Is our workflow robust enough for launching active learning?"

The term "robust" in traditional sense involves precision and accuracy, but in active learning, precision holds more weight. While accuracy pertains to the closeness of a measured value to the true value, the primary aim in active learning is to pinpoint the best candidate within the design space, rendering the relative order more crucial than the absolute value. Consequently, consistent deviation from the true value is tolerable as it doesn't

alter the relative performance ranking of candidates. Unlike accuracy, no compromise on precision would be acceptable other than random gaussian noise with fixed magnitude. In other terms, the results from the experiments have to be reproducible.

To tackle reproducibility issues, robotic arms are frequently utilized. While this aids in enhancing consistency to a degree, it's not a complete solution. Commonly, there's an expectation for automated experimental platforms to effortlessly yield consistent and reproducible outcomes, but this expectation is often met with disappointment in the early stages. To elucidate this point, we've prepared a simple website during the initial phase of constructing an automated platform in our lab (Ren, 2023c), showcasing the drop-casting process integral to the electrocatalyst optimization project. The goal here is to ensure uniform catalyst footprint across different samples. However, despite it's an automated process, noticeable variations in droplet position, diffusion distance, and wetting time are observed on each spot. Attaining the sought-after consistency took around three months, demanding careful control over numerous factors, including: (i) the wet-proofing layer of the carbon paper, (ii) the curvature and twist of the freshly-cut carbon paper strip, (iii) the anisotropy of the carbon fiber constituting the carbon paper, (iv) the gap between the pipette tip and the carbon paper, (v) the droplet dispensing rate. Neglecting any of these factors led to a lack of consistency, thereby hindering effective active learning optimization.

In brief conclusion, the underlying cause of irreproducibility is the presence of hidden variables, and these hidden variables originate from researchers' biases or ignorance towards the nature. For further insight, please refer to a summary of our observations during the application of active learning in real-world electrocatalysis experiments, encapsulated in a bulletin titled 'Criminals in the AI City' (Ren, 2023b). Currently, human researchers must identify these issues to enhance AI effectiveness in optimization tasks. However, we anticipate that in time, with the aid of multi-modal LLMs, AI will be sufficiently adept at uncovering these issues autonomously. Specifically, AI could compare all experimental logs, formulate scientific hypotheses to elucidate the causes of irreproducibility in data (Park et al., 2023), design and execute control experiments to validate these hypotheses, and manage the newly discovered hidden variables. In the long term, we envisage AI significantly advancing scientific research by generating hypotheses and steering the materials optimization process (Ren et al., 2023).

5. Conclusion

In summary, this study marks a significant advancement in materials science, particularly in the accelerated discovery of High-Entropy Alloys for electrocatalysis, through a novel integration of high-throughput robotics, active learning via BoTorch, and carbothermal shock fabrication. The addition of LLMs for voice-driven interactions further streamlines the experimental process. While these technologies show promise in optimizing green energy catalysts, the research emphasizes the critical need for addressing reproducibility challenges in real-world experiments. Success in this domain requires meticulous management of experimental variables to fully leverage the potential of AI.

Acknowledgments

We thank Ali Abdelhafiz, Wenhao Gao, Russell Scott In, David Liu, Yangjeong Park, Kai Pei, Zekun Ren, Hongbin Xu, Jialiang Zhao, Daniel Zheng for insightful discussions. We acknowledge support by DTRA (Award No. HDTRA1-20-2-0002) Interaction of Ionizing Radiation with Matter (IIRM) University Research Alliance (URA).

References

asweigart. Pyautogui. URL `https://pypi.org/project/PyAutoGUI/`.

Monya Baker. 1,500 scientists lift the lid on reproducibility. *Nature*, 533(7604), 2016.

Maximilian Balandat, Brian Karrer, Daniel R Jiang, Samuel Daulton, Benjamin Letham, Andrew Gordon Wilson, and Eytan Bakshy. BoTorch: A Framework for Efficient Monte-Carlo Bayesian Optimization. *arXiv*, 2019. doi: 10.48550/arxiv.1910.06403.

Benjamin Burger, Phillip M Maffettone, Vladimir V Gusev, Catherine M Aitchison, Yang Bai, Xiaoyan Wang, Xiaobo Li, Ben M Alston, Buyi Li, Rob Clowes, et al. A mobile robotic chemist. *Nature*, 583(7815):237–241, 2020.

Anubhav Jain, Shyue Ping Ong, Geoffroy Hautier, Wei Chen, William Davidson Richards, Stephen Dacek, Shreyas Cholia, Dan Gunter, David Skinner, Gerbrand Ceder, et al. Commentary: The materials project: A materials genome approach to accelerating materials innovation. *APL materials*, 1(1), 2013.

Benjamin P MacLeod, Fraser GL Parlane, Thomas D Morrissey, Florian Häse, Loïc M Roch, Kevan E Dettelbach, Raphaell Moreira, Lars PE Yunker, Michael B Rooney, Joseph R Deeth, et al. Self-driving laboratory for accelerated discovery of thin-film materials. *Science Advances*, 6(20):eaaz8867, 2020.

Yang Jeong Park, Daniel Kaplan, Zhichu Ren, Chia-Wei Hsu, Changhao Li, Haowei Xu, Sipei Li, and Ju Li. Can chatgpt be used to generate scientific hypotheses? *arXiv preprint arXiv:2304.12208*, 2023.

Jiayu Peng, James K Damewood, Jessica Karaguesian, Rafael Gómez-Bombarelli, and Yang Shao-Horn. Navigating multimetallic catalyst space with bayesian optimization. *Joule*, 5 (12):3069–3071, 2021.

Zhichu Ren. High-throughput robotic platform for electrocatalyst discovery, 2022. URL `https://youtu.be/iRauT95ECmo`.

Zhichu Ren. Crest - copilot for real-world experimental scientists, 2023a. URL `https://youtu.be/POPPVtGueb0`.

Zhichu Ren. Criminals in the ai city, 2023b. URL `http://li.mit.edu/cri`.

Zhichu Ren. Debugging a "simple" drop-casting process on carbon paper, 2023c. URL `http://li.mit.edu/automation`.

Zhichu Ren, Zekun Ren, Zhen Zhang, Tonio Buonassisi, and Ju Li. Autonomous experiments using active learning and ai. *Nature Reviews Materials*, pages 1–2, 2023.

RockChinQ. Callinggpt, 2023. URL `https://github.com/RockChinQ/CallingGPT`.

Ilya M Sobol. Global sensitivity indices for nonlinear mathematical models and their monte carlo estimates. *Mathematics and computers in simulation*, 55(1-3):271–280, 2001.

Eric Stach, Brian DeCost, A Gilad Kusne, Jason Hattrick-Simpers, Keith A Brown, Kristofer G Reyes, Joshua Schrier, Simon Billinge, Tonio Buonassisi, Ian Foster, et al. Autonomous experimentation systems for materials development: A community perspective. *Matter*, 4(9):2702–2726, 2021.

Helge S Stein and John M Gregoire. Progress and prospects for accelerating materials science with automated and autonomous workflows. *Chemical science*, 10(42):9640–9649, 2019.

Shijing Sun, Noor TP Hartono, Zekun D Ren, Felipe Oviedo, Antonio M Buscemi, Mariya Layurova, De Xin Chen, Tofunmi Ogunfunmi, Janak Thapa, Savitha Ramasamy, et al. Accelerated development of perovskite-inspired materials via high-throughput synthesis and machine-learning diagnosis. *Joule*, 3(6):1437–1451, 2019.

thanhsonng. Chatgpt-voice, 2023. URL `https://github.com/thanhsonng/chatgpt-voice`.

Yonggang Yao, Zhennan Huang, Pengfei Xie, Steven D Lacey, Rohit Jiji Jacob, Hua Xie, Fengjuan Chen, Anmin Nie, Tiancheng Pu, Miles Rehwoldt, et al. Carbothermal shock synthesis of high-entropy-alloy nanoparticles. *Science*, 359(6383):1489–1494, 2018.

Yonggang Yao, Qi Dong, Alexandra Brozena, Jian Luo, Jianwei Miao, Miaofang Chi, Chao Wang, Ioannis G Kevrekidis, Zhiyong Jason Ren, Jeffrey Greeley, et al. High-entropy nanoparticles: Synthesis-structure-property relationships and data-driven discovery. *Science*, 376(6589):eabn3103, 2022.