

# Position: Autonomous Scientific Discovery Needs Embodied Experimentation with Learnability

Yihang Chen<sup>\*1,2</sup> Zhongwei Yu<sup>\*3,2</sup> Yang Li<sup>\*†4</sup> Xingyu Lu<sup>5,2</sup> Qiang Li<sup>6</sup> Hongxia Yu<sup>7</sup> Jun Wang<sup>†1</sup>

## Abstract

Autonomous scientific discovery, particularly through self-driving laboratories (SDLs) integrated with large language models, has achieved breakthroughs in materials science, drug design, and beyond. However, current implementations exhibit a critical limitation: cognitive parameters remain immutable post-training, motor behaviours adhere to pre-programmed sequences, and experimental optimisation proceeds without leveraging institutional knowledge. Thus, we propose a formalised trinity of learnability comprising cognitive design, experimentation optimisation, and embodied execution, where an evolving cognitive core serves as the foundation for discovery. This architecture synergises the cognitive core with adaptive embodied skills to translate lab-specific variances into robustness while refining active experimental designs to prune explosive search spaces. This synergy enables SDLs to foster the emergence of high-efficiency discovery. **Taken together, this position paper argues that autonomous discovery requires learnability through embodied experimentation, not merely automation sophistication.** As a proof of concept, we developed a bespoke embodied laboratory prototype for the Electro-Fenton task. By leveraging historical memory and adaptive skills, our agent navigated a 200-million search space to achieve 99.7% efficiency in just 15 rounds.

## 1. Introduction

The automation of closed-loop scientific discovery has been a persistent ambition of the research community for over

<sup>1</sup>University College London <sup>2</sup>AI Lab, The Yangtze River Delta <sup>3</sup>The Hong Kong University of Science and Technology (Guangzhou) <sup>4</sup>Shanghai Jiao Tong University <sup>5</sup>Zhejiang University <sup>6</sup>Rein Agent <sup>7</sup>Nanjing Agriculture University. Correspondence to: Jun Wang <junwang@cs.ucl.ac.uk>, Yang Li <yang.li.cs@sjtu.edu.cn>.

Preprint. March 8, 2026.

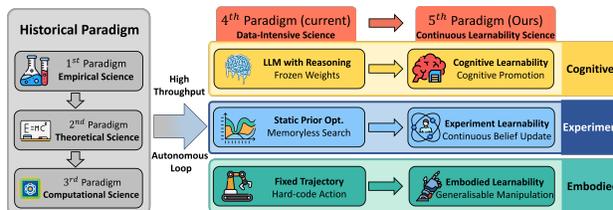


Figure 1. The evolution of scientific discovery paradigms. The historical progression from empirical observation to computational simulation has culminated in the current Fourth Paradigm - Data-Intensive Science, characterised by high-throughput yet static automation. We propose the Fifth Paradigm, Continuous Learnability Science, which transforms the system from a static autonomous platform into a discovery system with continuous learnability. This is achieved through the Trinity of Learnability, ensuring synchronised growth in Cognitive reasoning, Experimental optimisation, and Embodied execution, thereby enabling ontogenetic growth over the system’s lifetime.

half a century. From early mechanised synthesis (Merrifield, 1965; Merrifield et al., 1966; Zymark Corporation, 1982) to modern AI-driven platforms (King et al., 2004; 2009; Williams et al., 2015; Burger et al., 2020; MacLeod et al., 2020), the goal has remained constant to transcend the physical and cognitive limitations of the human researcher. This century-long pursuit has converged into the era of Artificial Intelligence (AI), which is fundamentally transforming scientific discovery and catalysing the emergence of AI for Science (AI4S) as a pivotal research frontier (Jumper et al., 2021; Yang et al., 2023; Noé et al., 2019). Within this ecosystem, Self-Driving Laboratories (SDLs) represent the vanguard of this movement. Unlike purely computational AI4S approaches (Noé et al., 2019; Merchant et al., 2023; Yang et al., 2023) that generate predictions from existing datasets, SDLs embody a closed-loop paradigm that bridges theoretical modelling and empirical experimental validation. By actively testing hypotheses through robotic experimentation in the physical world, SDLs enable an iterative cycle of hypothesis refinement and autonomous discovery (Burger et al., 2020).

As illustrated in Figure 1, the evolution of SDLs epitomises the Fourth Paradigm of Science (Data-Intensive Science) (Hey et al., 2009). Within this paradigm, SDLs have progressed through three distinct technological eras, each expanding the scope of this closed-loop automation. Early ef-

forts established foundational principles of systematic experimental design and pioneering automation, primarily serving as automation tools that executed predefined scripts (Merrifield, 1965; Merrifield et al., 1966; Zymark Corporation, 1982). The integration of advanced AI technologies has since catalysed a paradigm shift, elevating SDLs from narrowly scoped, protocol-driven platforms toward autonomous scientific agents (King et al., 2004; 2009; Williams et al., 2015; Steiner et al., 2019; MacLeod et al., 2020; Burger et al., 2020; Gupta et al., 2021). Modern platforms, such as Coscientist (Boiko et al., 2023), Organa (Darvish et al., 2025a), BioMARS (Zhao et al., 2025) and recent embodied intelligence-driven platforms (Li et al., 2025), leverage LLMs as reasoning engines to interpret complex instructions, integrate multimodal sensory data, and autonomously react to environmental contingencies. This leap has granted SDLs the ability to navigate the now with unprecedented sophistication. However, these advances mask a critical stagnation. While contemporary SDLs have become faster at executing experiments, they remain trapped in a paradigm of static cognition. In this state, the laboratory functions not as an intelligent researcher but as a high-speed execution engine operating on temporally fixed capabilities.

Building upon this critique, this position paper argues that **the fifth paradigm of scientific discovery requires learnability through embodied experimentation, not merely automation sophistication**. We conceptualise the modern laboratory through a three-layered framework comprising cognitive reasoning, embodied execution, and experimental optimisation, and we contend that the current stagnation in the field stems from a fundamental absence of learnability within each of these domains.

Cognitive learnability is essential to transform the laboratory from a static logic processor into an evolving strategist that internalises scientific laws and historical experimental insights to inform superior decision-making. By refining its own scientific heuristics over time, the system avoids remaining a perpetual novice that fails to upgrade its reasoning logic despite processing vast amounts of data. Similarly, embodied learnability must bridge the gap between digital commands and physical reality by enabling the platform to achieve robust generalisation across the complex environments of real-world laboratory variants. By utilising environmental feedback to perform dynamic adjustments, the system can navigate the inherent unpredictability of the physical world, where the purely scripted automation inevitably fails to achieve. Finally, experimental learnability is required to manage the prohibitive search spaces encountered in new discovery projects by building a persistent institutional memory. By internalising meta-knowledge across disparate campaigns, the laboratory can finally overcome the inefficiency of historical isolation and move beyond repetitive brute-force search toward the cumulative exper-

tise necessary for genuine scientific mastery.

By reframing these architectural gaps as a singular failure of learnability, we propose a transition toward an organism metaphor where the laboratory is viewed as an evolving scientific entity. This Trinity of Learnability functions as the core engine of this evolution, realised through three synergetic pillars: (i) Cognitive learnability is achieved by reflecting on memory-based case banks to internalise historical episodic experience and refine reasoning heuristics for future inquiries. (ii) Experimental learnability is attained through sequential Bayesian updates that transform noisy empirical observations into an evolving belief over vast hypothesis spaces. (iii) Embodied learnability is attained by acquiring and refining manipulation skills through reinforcement learning and high-fidelity simulation, allowing the system to compensate for physical noise and expand the repertoire of executable experiments; Through the co-evolution of these domains, the laboratory transcends its role as a static tool to become a unified ASD system capable of genuine cumulative growth.

To demonstrate the practical viability of this paradigm, we built a custom SDL specifically designed to instantiate the principles of systemic learnability. To demonstrate the efficacy of this paradigm, we engineered a custom-built SDL where the synergy of cognitive, embodied, and experimental learnability layers enabled the autonomous optimisation of Electro-Fenton wastewater degradation. Through this integrated loop, the system achieved a final transmittance of 99.7% in just 45 hours, effectively compressing a discovery process that traditionally requires two researchers and two sets of equipment over 30 days. This practical validation underscores the massive efficiency gains of our learnability framework can achieve rapid, data-driven mastery in complex real-world environments.

## 2. Related Work

Self-driving laboratories, defined as autonomous systems capable of formulating hypotheses, designing experiments, executing protocols, and iterating without human intervention, represent a grand challenge at the intersection of artificial intelligence, robotics, and scientific discovery. Over eight decades, SDLs have evolved from statistical experimental design frameworks through specialised robotic platforms to foundation model-driven systems, progressively expanding the scope of autonomous experimentation. Yet this evolution reveals a critical tension: increasing cognitive sophistication has come at the cost of physical grounding.

This section examines the development trajectory of SDL to understand both remarkable progress and persistent architectural limitations. Section 2.1 provides a historical perspective from the 1940s through the present, tracing

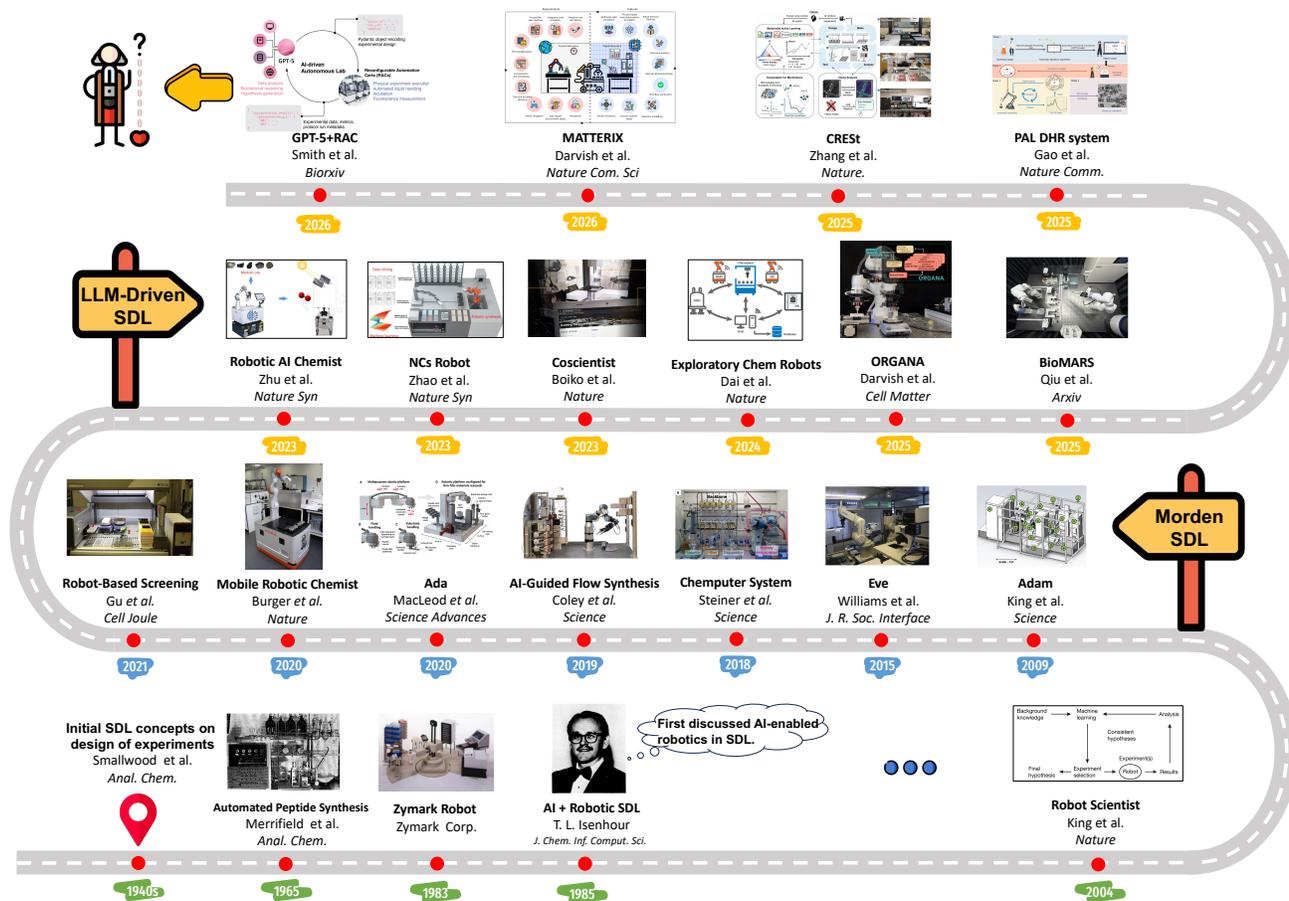


Figure 2. The Evolution of Self-Driving Laboratories. This timeline traces the development of autonomous experimentation systems from initial concepts in the 1940s to foundation model integration in recent years, spanning three eras: **Early Era** (1940s–2009, bottom) establishing foundational principles from design of experiments to closed-loop robot scientist; **Modern Era** (2009–2022, middle) deploying specialized robotic platforms across scientific domains; and **LLMs Era** (2022–present, top) integrating foundation models for high-level reasoning.

conceptual foundations and technological milestones. Section 2.2 analyses the robotics SDL roadmap (Fig. 2) across three eras: *Early Era* (1940s–2009) establishing foundations, *Modern Era* (2009–2022) deploying purpose-built platforms, and *LLM Era* (2022–present) integrating foundation models. Section 2.3 identifies open challenges that expose a fundamental paradox—current systems achieve linguistic reasoning while losing embodied understanding. This analysis motivates our paradigm shift toward embodied science agents that unify perception, cognition, and action within learned frameworks, as detailed in subsequent sections.

## 2.1. SDLs: A Historical Perspective

Early automation in laboratories began in the 1940s with parallel catalyst screening for the Haber-Bosch process (Topham, 1985). The 1960s saw Merrifield’s automated peptide synthesiser (Merrifield, 1965) and growing attention

to lab automation (Krasner, 1967). By the 1970s–1980s, closed-loop optimisation using simplex algorithms enabled autonomous tuning of chromatography and reaction conditions (Berridge, 1982; Matsuda et al., 1988), while Isenhour introduced the “Analytical Director” concept for AI-driven experiment planning (Isenhour, 1985). The 1990s brought combinatorial chemistry and high-throughput screening (Gordon et al., 1994), establishing feedback-driven workflows.

The 2000s marked a resurgence driven by advances in computing, machine learning, and robotics. Microfluidic reactors demonstrated closed-loop synthesis and optimization (Krishnadasan et al., 2007; Parrott et al., 2010; Kreutz et al., 2010; Adamo et al., 2016), while flow chemistry platforms (Coley et al., 2019; Christensen et al., 2021; Mehr et al., 2020) showcased autonomous reaction discovery. Landmark systems like *Adam* and *Eve* (King et al., 2009; Williams et al., 2015) demonstrated hypothesis-driven experimenta-

tion in genomics and drug screening. The late 2010s saw platforms autonomously discovering new reactions (Granda et al., 2018) and universal synthesis languages (Steiner et al., 2019). Materials discovery SDLs integrated Bayesian optimisation and neural networks for autonomous exploration (Raccuglia et al., 2016; Gongora et al., 2020; Häse et al., 2018; Roch et al., 2020; Seifrid et al., 2022).

Since 2020, SDLs have evolved into critical research infrastructure spanning chemistry, biology, and materials science (Lunt et al., 2024; Szymanski et al., 2023; Pyzer-Knapp et al., 2022). The emergence of LLMs, particularly following ChatGPT’s release in November 2022, catalysed a paradigm shift in how SDLs reason about and execute scientific tasks. LLM-based systems can now orchestrate tools via natural language, dramatically reducing barriers for complex workflows. Coscientist (Boiko et al., 2023) uses GPT-4 to plan experiments and control hardware via APIs, while ChemCrow (Bran et al., 2024) orchestrates synthesis planning and execution. AI4Sci Agent (Ding et al., 2025) proposes a unified framework for cross-disciplinary API coordination. Domain-specific systems address reaction prediction (Schwaller et al., 2021; Coley et al., 2017; Segler et al., 2018a), retrosynthesis (Segler et al., 2018b; Chen & Jung, 2021), materials design (Jablonka et al., 2024; Merchant et al., 2023; Kang et al., 2024), protein engineering (Hie et al., 2024; Madani et al., 2023), gene editing (Qu et al., 2025), and clinical applications (Singhal et al., 2023; Thirunavukarasu et al., 2023; Li et al., 2024; Goodell et al., 2025). Similar trends span chip design (Wu et al., 2024), mechanical prototyping (Ni & Buehler, 2024), geospatial analysis (Huang et al., 2024), neuroscience (Luo et al., 2025), and hazard response (Xie et al., 2025).

## 2.2. Robotics in Self-Driving Laboratories

**Early Era (1940s–2009).** As shown in the bottom of Fig. 2, the conceptual foundations for SDL were established decades before the technology existed to realise them. Smallwood’s 1947 work on the design of experiments in industrial research (Smallwood, 1947) laid crucial groundwork by formalising systematic approaches to experimental planning, introducing statistical methods that would later become essential for autonomous experiment design. This early emphasis on rational experimental design established principles that remain central to modern SDLs, particularly the optimisation of information gain from limited experimental resources.

The first significant robotic automation in chemistry emerged with Merrifield’s development of an automated peptide synthesiser in 1966 (Merrifield et al., 1966). This instrument demonstrated that complex sequential chemical operations could be reliably automated, synthesising peptides through programmable solid-phase chemistry. While not au-

tonomous in modern terms, Merrifield’s system proved that robotic platforms could execute intricate chemical protocols with precision and reproducibility, establishing automated synthesis as a viable paradigm. The work earned Merrifield the Nobel Prize in Chemistry in 1984, validating the importance of laboratory automation.

A pivotal conceptual milestone occurred in 1985 when Isenhour published “Robotics in the Laboratory” (Isenhour, 1985), explicitly articulating the potential for AI-enabled robotics in scientific discovery. This work directly addressed how intelligent robotic systems could transform laboratory science, presciently outlining many capabilities that would only become practical decades later. Isenhour’s vision encompassed not merely automated execution but intelligent experimental planning, recognising that true laboratory automation required computational reasoning integrated with robotic manipulation.

**Modern Era (2009–2022).** As shown in the middle of Fig. 2, the modern era of SDL development was characterised by increasingly sophisticated robotic platforms designed for specific scientific domains. The first genuine robotic scientist emerged in 2009 with King et al.’s groundbreaking Adam system (King et al., 2009). This platform autonomously generated hypotheses about yeast functional genomics, designed experiments to test those hypotheses, executed the experiments using robotic liquid handling, and analysed the results to form new hypotheses. Adam represented a qualitative leap beyond high-throughput screening, demonstrating that a machine could conduct the complete cycle of scientific inquiry. The system successfully identified the genes encoding orphan enzymes in yeast metabolism, making discoveries that were subsequently validated by human scientists.

King and colleagues extended this work in 2009 with Eve, a more sophisticated robotic scientist designed for drug discovery (Williams et al., 2015). Eve automated the screening of compounds against parasites responsible for neglected tropical diseases, autonomously selecting compounds for testing based on iterative learning. The system successfully identified several promising drug candidates, demonstrating that robotic scientists could conduct practically valuable research in applied domains. A major advance came in 2018 with Cronin’s group’s development of a modular robotic system driven by a chemical programming language (Steiner et al., 2019). This platform introduced a software abstraction layer that separated chemical synthesis logic from hardware implementation, enabling chemists to specify synthetic routes in high-level code that the robot automatically translated into physical operations. This architecture greatly enhanced flexibility and reproducibility, allowing the same synthesis code to run on different robotic configurations.

In 2019, Coley and colleagues demonstrated a robotic plat-

form for flow synthesis of organic compounds that integrated AI-based retro-synthetic planning with automated execution (Coley et al., 2019). The system could plan multi-step organic syntheses, optimise reaction conditions, and execute the protocols using continuous flow chemistry. This work highlighted the power of coupling modern AI methods for chemical retrosynthesis with robotic automation, enabling the synthesis of complex molecules with minimal human intervention. The platform successfully synthesised multiple pharmaceutical compounds, validating its practical utility.

A particularly ambitious project emerged from Cooper's laboratory in 2020: a mobile robotic chemist capable of navigating a standard chemistry laboratory and performing experiments using existing equipment (Burger et al., 2020). This system employed a mobile robot arm mounted on a wheeled base, along with computer vision for equipment recognition and localisation. The mobile chemist conducted photo-catalyst discovery experiments over several days of continuous operation, exploring a vast chemical space through active learning strategies. This demonstrated that SDLs need not require purpose-built equipment but could operate within conventional laboratory infrastructure, potentially democratising access to autonomous experimentation.

Also in 2020, MacLeod et al. reported an SDL for accelerated discovery of thin-film materials, focusing on photo-electrochemical water splitting catalysts (MacLeod et al., 2020). This system combined robotic thin-film deposition with automated characterisation and Bayesian optimisation to efficiently navigate composition spaces spanning multiple elements. The platform discovered several promising materials for solar fuel production, demonstrating the power of SDLs in materials science applications where traditional methods struggle with vast combinatorial spaces.

The application of SDLs to materials characterisation continued with robot-based high-throughput screening platforms for perovskite solar cells and other optoelectronic materials. Gu et al.'s 2020 work demonstrated automated screening of anti-solvents for lead halide perovskites (Gu et al., 2020), using robotic liquid handling combined with automated optical characterisation. These systems could process hundreds of conditions per day, far exceeding manual throughput while maintaining rigorous experimental control.

In 2022, Cronin's group presented a portable autonomous platform for universal chemical synthesis (Manzano et al., 2022), emphasising the miniaturisation and accessibility of SDL technology. This system, compact enough to transport between laboratories, could execute diverse chemical syntheses through a modular architecture supporting multiple reaction types. The portability aspect addressed a key limitation of many SDLs, their size and cost, potentially enabling

broader adoption across research institutions.

**LLM Era (2022–Present).** As shown at the top of Fig. 2, the integration of large language models and foundation models marked a transformative shift in SDL capabilities, enabling systems to leverage vast scientific knowledge encoded in pre-trained models. Boiko et al. demonstrated that large language models could autonomously conduct chemical research (Boiko et al., 2023), planning experiments, writing code to control laboratory instruments, and analysing results. The system successfully completed complex tasks, including optimising reaction conditions and discovering new catalysts, suggesting that LLMs could serve as the “cognitive engine” for SDLs, handling the reasoning and planning aspects while robots execute the physical operations.

A striking application emerged in 2023 when Zhu et al. reported an AI chemist that automated the synthesis of oxygen-producing catalysts from Martian meteorites (Zhu et al., 2024). This system analysed the composition of meteorite samples, autonomously designed catalyst formulations optimised for Mars conditions, and synthesised and tested the materials robotically. The work demonstrated SDL capabilities in an extreme application domain, extraterrestrial resource utilisation, showing that such systems could function with minimal human oversight in scenarios where human intervention is impossible.

Zhao et al.'s 2023 work on a robotic platform for colloidal nanocrystal synthesis (Zhao et al., 2023a) showcased how SDLs could tackle materials science challenges requiring precise control over nanoscale structures. The system autonomously optimised synthesis conditions for quantum dots and other nanocrystals, using robotic liquid handling combined with automated optical spectroscopy and machine learning for property prediction. This enabled rapid optimisation of nanocrystal properties for specific applications such as displays and photovoltaics.

In 2024, Dai and Cooper reported autonomous mobile robots for exploratory synthetic chemistry (Dai et al., 2024). These systems combined mobile manipulation with advanced AI for chemical exploration, autonomously investigating new reaction spaces and discovering novel chemical transformations. The robots could move between different experimental setups in a laboratory, combining multiple characterisation techniques to fully investigate chemical reactions. This work emphasised the growing sophistication of SDL decision-making, with systems capable of genuine exploratory behaviour rather than merely optimising within predefined spaces.

Recent developments in 2025 have focused on increasing the scope and accessibility of SDL technology. ORGANA, a robotic assistant for automated chemistry experimenta-

tion (Darvish et al., 2025a), provides sophisticated manipulation capabilities combined with computer vision for experiment monitoring. The system can conduct diverse chemical protocols while adapting to variations in experimental conditions, demonstrating robust performance across different chemical systems. Similarly, BioMARS introduced bimanual robotic manipulation for autonomous cell culture (Zhao et al., 2025), showing that SDL principles extend beyond chemistry to biological experimentation. These dual-arm systems can perform intricate manipulations required for cell biology protocols, potentially accelerating biological research.

The latest frontier of scientific discovery is defined by a transition from manual trial-and-error to cognitive, closed-loop automation. Exemplified by the work of Gao et al. (Gao et al., 2025), this evolution begins with complete end-to-end synthesis platforms for nanoparticles that integrate synthesis, purification, and characterisation into fully autonomous workflows. Building on this physical autonomy, Zhang et al. (Zhang et al., 2025) introduced the CRES platform, which leverages multimodal vision-language models to accelerate catalyst discovery by synthesising prior knowledge from literature and microscopic imagery. To mitigate the costs associated with such physical exploration, Darvish et al. (Darvish et al., 2025b) developed the MATTERIX framework, a high-fidelity digital twin that utilises a GPU-accelerated semantic engine to bridge the "sim-to-real" gap in chemical kinetics. Finally, moving into the realm of complex biological logic, Smith et al. (Smith et al., 2026) demonstrated the potential of GPT-5 driven agentic AI to manage reconfigurable automation arrays, achieving unprecedented efficiency in cell-free protein synthesis. Together, these advancements represent the maturation of SDL technology into self-evolving tools capable of conducting independent, publishable research across diverse scientific domains.

### 2.3. Open Challenges and Opportunities

Despite remarkable advances across three technological eras, current foundation model-era SDL systems exhibit a fundamental architectural inadequacy rooted in the absence of embodied intelligence: they treat laboratory experimentation as abstract symbol manipulation rather than embodied physical practice. Human scientists develop experimental expertise via embodied interaction, gaining sensorimotor experience that certain colour changes indicate reaction completion, specific stirring patterns improve mixing efficiency, or reagent viscosity requires adjusted pipetting speeds. This tacit knowledge, acquired through years of physical practice, enables adaptive problem-solving when experiments deviate from expectations—recognising when a colour change signals an interesting side reaction versus contamination, or when equipment malfunction can be circumvented through alter-

native procedures. In contrast, current SDL architectures maintain strict separation between disembodied foundation models (handling reasoning and planning) and robotic systems (executing pre-programmed instructions), fundamentally preventing the development of such grounded understanding. This architectural choice manifests in three critical limitations that constrain autonomous scientific capability:

**(1) Lack of Physical Grounding.** Current systems can perform laboratory operations such as mixing reagents, heating solutions, or observing colour changes, but they lack an understanding of the underlying physical principles. LLMs generate experimental plans based on statistical patterns in scientific literature rather than causal models grounded in physical experience. This leads to plans that are syntactically correct but physically infeasible: suggesting stirring speeds that cause splashing, heating rates that decompose intermediates, or reagent additions that violate solubility constraints. The systems lack perceptual grounding to recognise visual indicators (precipitation, phase separation, colour intensity) that guide human experimentalists in real-time decision-making.

**(2) Rigid Separation Preventing Adaptive Execution.** Existing architectures enforce strict boundaries between cognitive modules (foundation models for planning) and execution modules (robotic systems). This separation prevents real-time adaptation when unexpected situations arise. The planning module, lacking sensory perception, cannot observe that a reagent has higher viscosity than expected and adjust pipetting accordingly. The execution module, lacking reasoning capabilities, cannot recognise that an unexpected colour change might indicate an interesting side reaction worth investigating. When experiments deviate from predicted trajectories, such as reagent spills, equipment malfunctions, or unexpected reaction kinetics, systems typically fail catastrophically or require human intervention, as they are unable to improvise solutions through integrated perception, reasoning, and action.

**(3) Absence of Sensorimotor Learning and Skill Improvement.** Foundation model-era systems do not develop experimental intuition through repeated physical interaction. While modern SDLs can optimise experimental parameters (temperatures, concentrations, reaction times) through active learning algorithms, they do not improve their fundamental experimental skills through practice. A system executing its thousandth titration performs no better than its first because motor skills (precise liquid handling, equipment manipulation) and perceptual judgments (recognising endpoint indicators, detecting equipment issues) remain fixed. Current architectures lack mechanisms for closed-loop learning where physical execution outcomes inform skill refinement, preventing the development of tacit knowledge that human scientists accumulate through hands-on

experience.

**Summary.** These limitations are not isolated engineering challenges but symptoms of a deeper architectural mismatch. Our historical analysis reveals a paradoxical trajectory: SDL evolution from fixed-protocol automation (1940s–2009) through purpose-built platforms (2009–2022) to foundation model integration (2022–present) has increased cognitive sophistication while decreasing physical grounding. Systems gain linguistic reasoning and knowledge retrieval at the cost of embodied understanding, treating experimentation as symbol manipulation rather than physical practice. This paradigm, treating language models as disembodied planning oracles that generate instructions for separate robotic executors, fundamentally contradicts how human scientists develop expertise. Human experimental competence emerges through the inseparable integration of conceptual understanding, perceptual skills, and motor competence developed over years of hands-on laboratory practice. A chemist’s ability to “know” when a reaction is complete integrates theoretical knowledge (expected reaction kinetics), perceptual experience (recognising characteristic colour changes), and motor memory (stirring patterns that ensure complete mixing). Current SDL architectures, by separating these capabilities across disconnected modules, cannot develop such integrated expertise.

Addressing these limitations requires a paradigm shift toward *embodied science agents* that unify perception, cognition, and action within integrated frameworks. Rather than treating foundation models as disembodied planners, embodied science agents must cultivate embodied learnability by developing a grounded understanding through direct sensorimotor interaction with laboratory environments, continuously learning physical constraints, acquiring perceptual skills, and refining motor competence through repeated practice, much as human scientists do.

### 3. Why Learnability Matters

This work examines the historical progression of SDLs across three distinct technological eras while focusing on the architectural metamorphosis from early script-bound automation to contemporary reasoning-driven autonomous agents. By evaluating these developments, we pinpoint the systemic learnability gaps within current frameworks that hinder cognitive and physical evolution, thereby establishing the necessity for a transition toward the Fifth Paradigm of autonomous scientific discovery. **The Critical Gap: From Automation to Discovery** The earliest landmark progress in modern SDLs originated from the first implementation of an end-to-end autonomous discovery loop within a physical system. In 2004, the team led by Ross D. King (King et al., 2004) demonstrated a robot scientist system in the field of functional genomics, proving that the complete sci-

entific methodology, comprising hypothesis generation, experimental design and execution, could be formalised and executed autonomously by machines in a closed-loop fashion. Building upon this, the Adam system (King et al., 2009) proposed in 2009 further systematised these principles by explicitly encoding the scientific reasoning cycle while emphasising machine-readable experimental representations and the standardised recording of data and processes. These early efforts established the foundational technical architecture for robotic scientists through knowledge representation and iterative learning modules.

The evolution of robotic discovery has naturally transitioned from the formalisation of scientific logic toward the autonomous navigation of complex experimental spaces. This paradigm shift is exemplified by landmark systems such as the Chemputer (Steiner et al., 2019), which utilises a standardised chemical programming language for modular synthesis, and the mobile robotic chemist (Burger et al., 2020), which achieved a breakthrough in physical autonomy by independently identifying superior photocatalysts through hundreds of experiments. Despite these remarkable achievements in throughput, these platforms remain anchored to a fixed capability boundary and exhibit a form of institutional amnesia. Because their underlying reasoning weights and motor policies are immutable post-deployment, every new task begins as a cold start, failing to internalise the empirical hand-feel, the site-specific nuances of sensor noise and hardware tolerances required for scientific mastery. Without the capacity for ontogenetic growth, the laboratory remains a sophisticated but longitudinally static instrument rather than an evolving scientific organism.

Building upon these foundations, the integration of LLMs and tool-augmented agent architectures has introduced a sophisticated orchestration layer to autonomous discovery. Modern frameworks like Organa (Darvish et al., 2025a) and SciToolAgent (Ding et al., 2025) leverage foundation models as reasoning engines to translate abstract research goals into complex experimental workflows while coordinating diverse physical instruments. While this leap provides unprecedented flexibility and cross-domain generalisation, it primarily enhances the laboratory’s sophistication rather than its agency. These systems still exhibit static intelligence; because their underlying cognitive parameters remain immutable post-deployment, they cannot internalise empirical feedback to mature their scientific intuition. Consequently, despite their advanced reasoning capabilities, these laboratories serve merely as transient executors that reset to a baseline state after every task, lacking the historical continuity required for true mastery.

In summary, the transition from data-intensive automation to intelligence-intensive evolution marks the dawn of the Fifth Paradigm of Science: continuous learnability. Unlike the

Fourth Paradigm, which scales experimentation, this new paradigm demands learnability to bridge the gap between static tools and evolving agents. It necessitates a transformation across three fundamental dimensions: cognitive learnability to reflect on reasoning failures and refine scientific intuition from historical memory; Experiment learnability to continuously update beliefs and evolve search policies within vast hypothesis spaces; and embodied learnability to master the complex physical environment by acquiring adaptive manipulation skills. By integrating this tripartite, learnability can a SDL can transcend the institutional amnesia of the Fourth Paradigm and achieve the ontogenetic growth of a truly expert scientific discovery organism.

#### 4. Towards the Fifth Paradigm: A Unified Framework

To realise the vision of the Fifth Paradigm, SDLs must move beyond the mere computerisation of prescribed patterns and instead embed lifelong learnability into their frameworks. Unlike the traditional training-and-testing paradigm, lifelong learnability emphasises that experience accumulates continuously as the system operates, rather than relying on static snapshots of data. However, learnability in existing SDLs remains fundamentally limited, often confined to local modules or fragmented into independent loops that lack collective evolution. Crucially, this evolution must be holistic. Guided by the Cannikin Law, we posit that a scientific agent is only as effective as its least learnable component, as a failure to adapt in execution or strategy will inevitably bottleneck even the most advanced reasoning engine. To eliminate these bottlenecks, we delineate three foundational pillars where the system must exhibit synchronised growth, as illustrated in Figure 3. We refer to these as the Trinity of Learnability.

##### 4.1. The Trinity of Learnability

In this section, we delineate the three foundational pillars where an autonomous system must exhibit learnability to achieve true scientific autonomy, as illustrated in Figure 3. We refer to these as *the trinity of learnability*: cognitive reasoning, experimentation, and embodied manipulation. While we primarily focus on SDLs operating under the H-D paradigm—the standard in most modern laboratories—the core concepts generalise to any scientific method requiring empirical evidence.

###### 4.1.1. COGNITIVE LEARNABILITY

Cognitive reasoning represents the highest level of the discovery loop, where the system translates a high-level user query into the constituent elements of a scientific workflow, typically the H-D methodology in the context of SDLs.

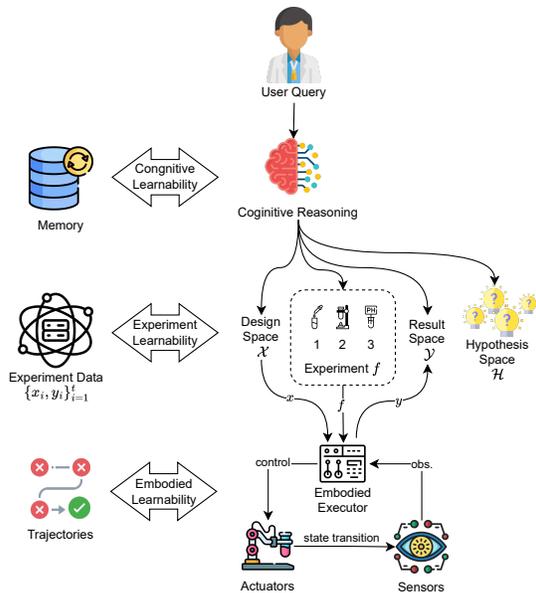


Figure 3. The trinity of learnable aspects in autonomous scientific discovery. The cognitive layer reasons the task hypotheses and parametric experiment procedures from the user query; the experimentation layer designs experiments to test these hypotheses; and the embodied layer executes experiments by interacting with the physical environment.

Formally, given a query and laboratory context, cognitive reasoning formalises the task into a representation:

$$\text{query} \mapsto (\mathcal{H}, \mathcal{X}, \mathcal{Y}, f)$$

Here,  $\mathcal{H}$  denotes the hypothesis space of unknown components requiring empirical evidence, and  $f : \mathcal{X} \rightarrow \mathcal{Y}$  represents the executable experimental procedure used to test these hypotheses, where  $\mathcal{X}$  is the space of experimental designs and  $\mathcal{Y}$  is the space of observable results. Effective cognitive reasoning must ensure that both the hypothesis space and the experimental design remain relevant to the user’s query. Specifically, experiments  $f$  across various designs should effectively falsify or support specific hypotheses in  $\mathcal{H}$ , such that the true answer to the query is determined by which hypothesis holds.

With the integration of foundation models, cognitive reasoning can now be automated. By prompting LLMs, an SDL can generate formatted task descriptions where elements  $(\mathcal{H}, \mathcal{X}, \mathcal{Y}, f)$  are expressed in code or a prescribed grammar. However, learnability at this level is often absent in contemporary SDLs, which frequently rely on static templates or hard-coded prompts.

To advance, an SDL must exhibit *metacognitive learnability*—the ability to reflect on the success or failure of its previous reasoning chains. A direct approach is to fine-tune the large-scale foundation models, which, however, requires

massive computational resources and data. A more tractable approach is *learning how to reflect from memory* (Zhou et al., 2025; Wang, 2025). Formally, let  $\mathbf{a}_i = (\mathcal{H}_i, \mathcal{X}_i, \mathcal{Y}_i, f_i)$  be the tuple representing the cognitive reasoning that the SDL has produced for the  $i$ -th query  $q_i$ , and let  $r_i$  be the measurement of discovery performance (e.g., information gain over the hypothesis space, or criteria of experiment results). The SDL then maintains a memory of all encountered queries:  $\mathcal{M} = \{(q_i, \mathbf{a}_i, r_i)\}_{i=1}^m$ . Equipped with such a case bank, the SDL learns a policy  $\pi(\mathbf{a}|q, \mathcal{M})$  to retrieve relevant methodological patterns and adjust its planning for future discoveries, effectively evolving its “intuition” for scientific inquiry.

#### 4.1.2. EXPERIMENT LEARNABILITY

Once a scientific procedure  $f$  is established, the system must determine which specific designs  $x$ , aiming to disclose the key information regarding hypothesis  $h$  and therefore contributing to the user’s query. Statistical experiment design (Box et al., 2005) provides the principled tools to address this. To enable learnability, experiment design is conducted sequentially, allowing new designs to be informed by previous observations.

Because the hypothesis space  $\mathcal{H}$  in modern science is often enormous or infinite, testing hypotheses individually is usually intractable. Furthermore, due to manipulation errors and uncontrollable factors,  $f$  often functions as a noisy black box where the outcome  $y$  is perturbed from the true hypothesis. Consequently, we propose a shift from deterministic testing to a *Bayesian extension* of the scientific method. In this framework, the SDL maintains a belief over the hypothesis space  $\mathcal{H}$ , which is updated continuously as empirical evidence  $y$  is collected.

By incorporating Bayesian principles into H-D discovery, we enable Bayesian learning for sequential experiment design. For instance, *Bayesian experimental design* (BED) (Rainforth et al., 2023) can be applied to queries seeking to explain unknown mechanisms. Conversely, many queries are objective-directed, aiming to determine how to act within an environment with unknown mechanisms; in such cases, *Bayesian optimisation* (BO) serves as a principled tool that extends BED by optimising designs with respect to black-box objectives. This shift from brute-force enumeration to strategic navigation enables the agent to locate global optima with orders of magnitude fewer experiments.

#### 4.1.3. EMBODIED LEARNABILITY

Scientific discovery is rarely a purely cerebral exercise; it requires physical interaction with the environment. Embodied learnability ensures that the experimental procedure  $f$  is executed reliably by low-level actuators. In a physical laboratory,  $f$  is not merely a mathematical function but a

sequence of physical interactions, such as pipetting, heating, or titrating, that are subject to mechanical tolerances and environmental noise.

Existing SDLs often suffer from “brittleness” due to rigid control programs that fail when spatial configurations change, or hardware degrades. A learnable SDL must instead acquire and refine manipulation skills through experience. This is typically achieved through Reinforcement Learning (RL), where an agent learns optimal control policies via trial and error. To mitigate the risks and costs of real-world failures, high-fidelity physical simulators (e.g., MuJoCo or Isaac Lab) allow the SDL to “practice” and acquire robust skills before deployment (Mittal et al., 2023; Todorov et al., 2012). As the system operates, it learns to compensate for calibration errors and masters new instruments, effectively expanding the repertoire of experiments the cognitive layer can feasibly propose, while ensuring robust execution even in the presence of modelled hardware noise.

## 4.2. Synergistic Learnability

The true power of the Trinity lies not in the isolated optimisation of its components, but in their deep integration. Rather than functioning as separate modules, cognitive reasoning, experimentation, and embodied manipulation must operate as interlocking gears of a single engine. In this unified framework, the constraints of one layer become the optimisation objectives of another, driving the system toward a level of autonomy that no single component could achieve alone.

Together, these components form a cohesive closed-loop system of self-evolution where progress in one domain lowers the barrier for advancement in others. From a top-down perspective, improved cognitive reasoning provides accurate prior beliefs that expedite the convergence of Bayesian experimentation, while feedback from embodied control ensures that high-level planning remains physically tractable. Conversely, as the embodied layer masters precise manipulation skills, it reduces the evaluation noise of experimental results ( $y$ ), thereby sharpening the strategic search. This physical mastery effectively expands the system’s operational envelope, allowing the strategic layer to explore regions of the hypothesis space that were previously inaccessible due to mechanical limitations. Finally, the high-quality data yielded by this efficient process is internalised into institutional memory, providing the rich empirical experience necessary for the cognitive layer to distil generalizable principles, thus completing the cycle of lifelong learning.

This mutual reinforcement directly addresses the Cannikin Law bottleneck introduced at the beginning of this section. It ensures that improvements in reasoning are not stifled by a failure to adapt in execution or strategy. By coupling

these layers, the architecture transitions from a “tabula rasa” novice into an evolving expert agent. Ultimately, this synergy effectively decouples experimental complexity from the time required for discovery, allowing every experimental cycle to contribute to a cumulative reservoir of scientific expertise.

## 5. Case Study: Proof of Concept

To bridge the gap between theoretical frameworks and physical reality, we build a custom SDL prototype and present a comprehensive case study centred on the Electro-Fenton degradation of Safranin T, a common industrial dye used to simulate the removal of organic pollutants from wastewater. This scenario represents a quintessential high-dimensional, physically-grounded challenge characterised by the non-linear interplay of electrochemical kinetics and the stochastic nature of robotic manipulation.

While conventional self-driving laboratories would treat this task as an isolated optimisation problem, our implementation serves as a validation of the Learnable Embodied Experimentation paradigm. By integrating a memory-augmented cognitive layer, a visuomotor-driven embodied layer, and a data-driven experimental layer, the system demonstrates how an agent with trinity learnability can transcend the static intelligence of previous eras. In the following sections, we detail how the system leverages its initial cognitive priors derived from mock methodologies, refines its embodied hand-feel through VLA-based manipulation, and achieves a 99.7% degradation efficiency through the continuous modelling of experimental experience by BO.

### 5.1. Prototype: Embodied Science Agent

As a proof of concept that operationalises this Trinity of Learnability, we introduce the Embodied Science Agent (ESA), a hybrid cognitive-embodied framework designed for autonomous self-driving laboratories, as illustrated in Figure 4. To realise the synergistic closed-loop evolution discussed previously, ESA tightly couples digital intelligence with physical experimentation. At the highest level, the Digital Agent acts as the cognitive cerebrum, featuring a Principal Investigator (PI) Agent that translates user intents into graph-based workflows. Directly addressing cognitive learnability, this agent leverages deep research and an episodic memory module, where past experimental plans and empirical information are continuously written and accumulated, to reflexively refine its methodological intuition over time. Once a cognitive plan is formulated, the Embodied Unified Interface (EUI) serves as the critical coordination bridge, leveraging the MCP protocol to securely dispatch these tasks without relying on rigid digital-twin abstractions. This ensures seamless integration with the Embodied Agent, the system’s cerebellum, which physi-

cally grounds the abstract plans through robotic systems and scientific instruments. Fulfilling the mandate of embodied learnability, this execution layer translates instructions into precise real-world manipulations while managing physical environmental noise. Following execution, the loop is closed by feeding the empirical observations back to a Bayesian Optimiser; embodying experiment learnability, this optimiser sequentially explores the parameter space and proposes new designs for the PI Agent to replan. This iterative, mutually reinforcing cycle continues until termination criteria are met, ultimately yielding a comprehensive scientific report and demonstrating how cognitive reasoning, Bayesian experimentation, and physical execution can function as interlocking gears to achieve open-ended scientific discovery.

### 5.2. The SDL Prototype with Learnability

The physical realisation of our framework is an integrated, multi-modal laboratory environment (Figure 5) designed to serve as a high-fidelity interface between digital reasoning and physical experimentation. The platform orchestrates a suite of analytical and executive hardware—including magnetic stirrers, programmable power supplies, air and syringe pumps, and spectrophotometers—into a unified orchestration layer. This layer ensures that every hardware state and experimental variable is digitised and timestamped, providing the continuous data stream necessary to build the system’s institutional memory.

For legacy instruments lacking digital APIs, such as the pH meter and the electrochemical workstation, we implement a closed-loop vision-based interface. This subsystem utilises high-resolution cameras and optical character recognition (OCR) to bridge the gap between physical readouts and digital reasoning, effectively transforming analogue displays into real-time telemetry. Physical labour is performed by a dual-arm collaborative robotic system equipped with custom end-effectors. Unlike single-arm systems restricted to rigid linear paths, this anthropomorphic configuration allows for complex, multi-step tasks such as positioning pH probes, electrode handling, and simultaneous reagent titration. This dexterity is critical for embodied learnability, as it provides the system with the mechanical degrees of freedom required to refine its motor policies based on environmental feedback. Detailed specifications regarding the dual-arm robotics and the complete inventory of experimental equipment are provided in Appendix A.

### 5.3. Implementation of Learnability

The learnability of this prototype is realised through the tight coupling of the three layers, transforming static automation into an evolving scientific agent. For the memory-driven methodology within the cognitive layer, the system is ini-

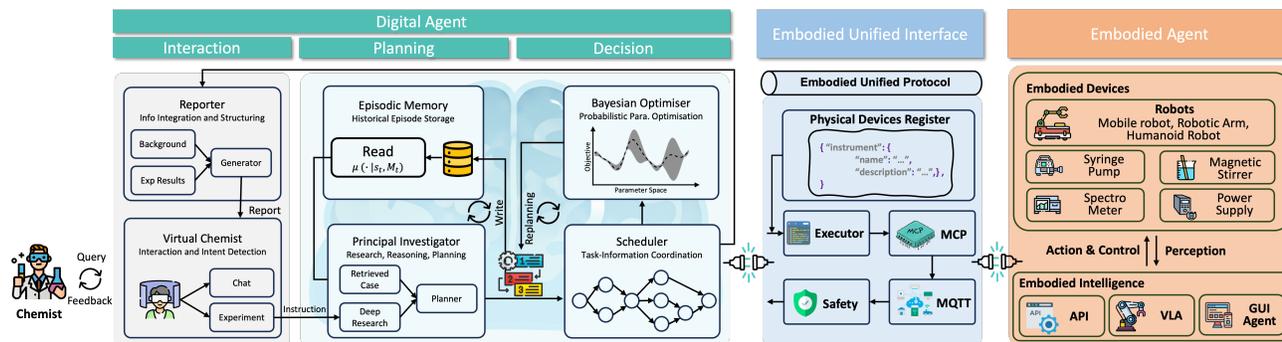


Figure 4. The Architecture of the Embodied Science Agent. ESA operationalises the *Trinity of Learnability* by establishing a hybrid cognitive–embodied architecture that tightly couples digital intelligence with physical experimentation. The cognitive core, the **Digital Agent**, drives *cognitive and experiment learnability* through open-ended evolution, integrating online knowledge, Bayesian experimentation, and accumulated past experience to continuously refine its reasoning and planning capabilities. The **Embodied Agent** orchestrates diverse robotic and scientific instruments to realise *embodied learnability*, enabling adaptive, reliable physical execution that robustly manages environmental noise across dynamic setups. The **Embodied Unified Interface** supports this integration by securely connecting digital intelligence with heterogeneous laboratory devices via standardised, low-latency communication protocols. Together, these components form a synergistic closed loop of self-evolution, where continuous feedback between reasoning, experiment design, and embodied execution embeds lifelong learnability into the autonomous scientific discovery process.

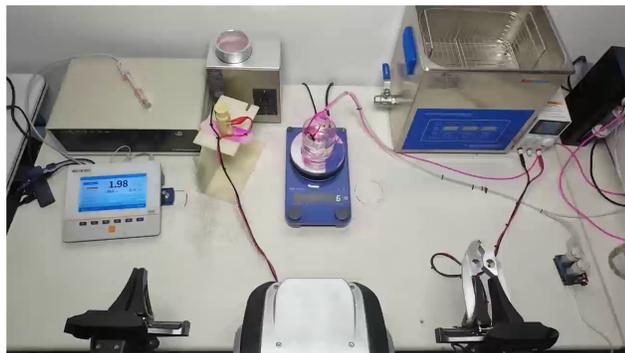


Figure 5. The physical realisation of the custom SDL prototype. The integrated environment features a dual-arm robot orchestrating a comprehensive suite of laboratory hardware. The experimental setup comprises executive devices (magnetic stirrers, programmable power supplies, air and syringe pumps) and analytical instruments (spectrophotometers, electrochemical workstations, and pH meters).

tialised with structured mock chemical methodologies that define the basic protocols for acidification and treatment. Unlike previous paradigms that treat these steps as fixed code, our framework treats them as evolvable templates. The Cognitive Layer “experiences” the instrument interfaces through vision-based screen understanding, extracting data to internalise environmental feedback. Crucially, the experimental plans and observed outcomes are written into a long-term memory, allowing the agent to refine its internal models of the laboratory’s specific physics for future reuse.

The experiment layer manages a vast search space comprising over 200 million unique experimental configurations, including pH, voltage, aeration, and electrolyte concentrations. Learnability here is manifested as the continuous refinement of a BO surrogate model. Each experimental

Table 1. The discretised search space for Electro-Fenton parameter optimisation.

Parameter	Symbol	Range / Set	Step Size	Unit
Acidity	pH	{2.0, 2.5, 3.0, 4.0, 5.0}	(Set)	-
Applied Voltage	$U$	1.00 – 10.00	0.01	V
Aeration Intensity	PWM	50 – 255	1	-
Stirring Rate	$\omega$	100 – 300	10	rpm
$\text{Na}_2\text{SO}_4$ Conc.	$C_{\text{Na}_2\text{SO}_4}$	25 – 75	5	mmol L <sup>-1</sup>

round updates the system’s strategic intuition, allowing the agent to learn the intricate interplay between parameters, such as the relationship between voltage and pH stability, and identify the global optimum through experience. The implementation details of BO for this scenario, through Gaussian Processes (GPs) and the upper confidence bound (UCB) is in Appendix B.

In the Embodied Layer, the transition from rigid scripts to adaptive policies is achieved through training specialised visuomotor policies. The robotic arms do not rely on hard-coded coordinates; instead, they utilise vision-language-action (VLA) strategies to handle objects with varying positions and orientations. As demonstrated in the generalisation trials, the agent has learned the physics of the workspace, enabling it to successfully grasp electrode racks and insert probes even when the apparatus is displaced. This ability to acquire spatial intuition ensures that the laboratory remains operational despite physical stochasticity.

#### 5.4. Experimental Configuration

In this case study, Safranin T was chosen as the target pollutant to evaluate the Electro-Fenton degradation performance of the fabricated iron–sludge composite electrode. The optimisation of operational parameters was carried out using

this dye as the model compound, allowing a systematic assessment of the electrode’s degradation efficiency. Based on their critical influence on hydroxyl radical generation and pollutant degradation, this work focuses on five key parameters:

- **Acidity:** controls the  $\text{Fe}^{2+}/\text{Fe}^{3+}$  cycle and the stability of hydroxyl radicals.
- **Voltage:** determines the potential applied across the electrodes, affecting  $\text{H}_2\text{O}_2$  generation and the rate of radical formation.
- **Aeration rate:** determines the supply of  $\text{O}_2$ , which influences  $\text{H}_2\text{O}_2$  generation.
- **Stirring rate:** governs mass transfer and the distribution of reactive species within the solution.
- **$\text{Na}_2\text{SO}_4$  concentration:** acts as the supporting electrolyte, affecting conductivity and the electrochemical environment.

Crucially, these parameters do not operate in isolation; their intricate interplay gives rise to significant interaction effects. For instance, the influence of the applied voltage is closely intertwined with pH: high voltage at low pH can accelerate electrode corrosion and intensify parasitic hydrogen evolution, thereby reducing current efficiency. Similarly, pH interacts with the concentration of iron species; at higher pH values, iron ions form precipitates, whereas at excessively low pH, hydroxyl radical generation is suppressed. These combinatorial interactions create a vast search space containing millions of potential configurations, rendering conventional one-factor-at-a-time (OFAT) methods impractical due to prohibitive resource demands. Within the framework of learnability, the Cognitive Layer initially constrains this vast theoretical space into a tractable set of discretised bounds based on historical priors. The Experiment Layer then explores this space via the BO module to maximise the final solution transmittance (ranging from 0% to 100%). The specific parameter bounds and step sizes for the embodied execution are detailed in Table 1.

### 5.5. Empirical Evidence

The optimisation trajectory (Figure 6) validates this learnable architecture. Starting from a baseline transmittance of approximately 82%, the process began with five rounds of structured exploration. After a plateau phase where the system refined its understanding of parameter interactions, it achieved a pronounced breakthrough. By the final iteration (Round 15), the system identified a globally optimal configuration that reached a transmittance of 99.7%, representing near-complete degradation of the target pollutant. This leap demonstrates that a learnable agent can effectively decouple experimental complexity from time, evolving from an intelligent novice into an expert scientific operator.

The efficiency gains are substantial. Navigating the vast

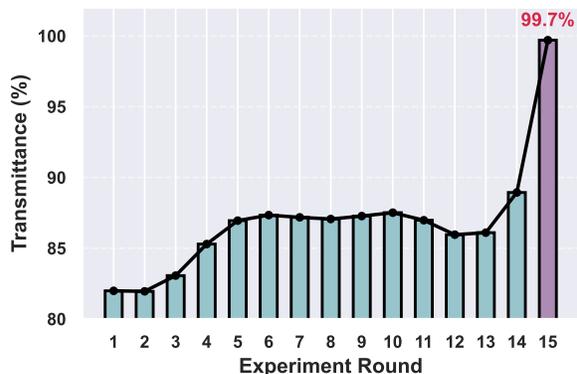


Figure 6. Transmittance results across 15 experimental rounds. The chart tracks solution transmittance across 15 experimental rounds. Starting from a baseline of  $\approx 82\%$ , the agent progresses through exploration (Rounds 1–5) and refinement phases (Rounds 6–13) before identifying the global optimum in Round 15, achieving a final transmittance of **99.7%**.

search space to locate this global optimum would impose a prohibitive toll on human resources; a traditional manual campaign of this scale, requiring sufficient replicates for statistical significance, would demand two researchers working for approximately 30 days. In stark contrast, our prototype achieved this expert-level proficiency using a single set of equipment within just 45 hours.

Crucially, this rapid convergence is not attributable to the optimisation algorithm alone, but rather to the synergistic interplay of the Trinity of Learnability. The reliability of the optimisation results is fundamentally grounded in the Embodied Layer, where adaptive motor skills ensured precise manipulation and stable sensing, thereby minimising physical execution noise and providing the high-fidelity feedback necessary for valid modelling. Simultaneously, the Cognitive Layer leveraged historical memory to define a tractable initial search space, preventing the system from wasting iterations on chemically implausible regions. It was on this foundation of high-quality data and valid search boundaries that the Experiment Layer could efficiently update its beliefs to navigate out of local optima. Thus, the result proves that autonomous scientific mastery is an emergent property arising from the co-evolution of cognitive reasoning, experimental optimisation, and embodied execution.

## 6. Alternative views

### 6.1. Pre-training vs. Continuous Adaptation

**Position** The mainstream paradigm for AI deployment in the current era is pre-training and fine-tuning. Proponents argue that sufficiently large foundation models, trained on vast scientific literature and simulations, provide emergent zero-shot generalisation capable of handling most experimental

variants. From an engineering perspective, this centralised approach ensures superior stability, lower maintenance, and cross-laboratory standardisation compared to complex, self-evolving local architectures.

**Rebuttal** We acknowledge that foundation models provide an unparalleled “scientific cerebrum” of reasoning. However, we contend that site-specific learnability and pre-trained priors are orthogonal dimensions of intelligence. While pre-training and fine-tuning excels at capturing global scientific consensus, it inherently manifests as static intelligence once deployed. It cannot account for site-specific physics, such as idiosyncratic sensor noise, reagent batch variations, or non-linear mechanical tolerances, that are unique to a particular laboratory. Our framework treats learnability not as a replacement for foundation models, but as an interface that allows static priors to mature into dynamic expertise. By internalising real-world feedback into a long-term memory, the system bridges the reality gap that frozen, pre-trained weights alone cannot traverse. Furthermore, as argued by several relevant works (Ding & Li, 2025; Maffettone et al., 2023), generative AI alone lacks the fundamental creativity required for zero-shot scientific discovery from scratch. This reinforces our argument that an embodied, evolving agent is necessary to complement static foundation models, bridging the gap between hallucination and reality.

## 6.2. Deterministic Precision vs. Adaptive Uncertainty

**Position** A compelling counter-argument is that research productivity is primarily a function of automation sophistication—the peak physical throughput and deterministic precision of the hardware. Proponents of this High-Throughput First paradigm argue that resources are most effectively spent on expanding the parallelisation and speed of robotic platforms, such as those pioneered by Steiner et al. (Steiner et al., 2019) or Burger et al. (Burger et al., 2020). From this perspective, an SDL should be a high-fidelity repeater that minimises variance through rigid, instruction-following pipelines. They contend that introducing learnability into the loop, especially through black-box systems, risks compromising the interpretability and reproducibility of scientific data, as an evolving organism might introduce non-linear biases that are difficult to calibrate.

**Rebuttal** While we acknowledge the importance of precision, we argue that the sophistication-only approach leads to a Fixed Capability Boundary. In high-dimensional spaces, like our 200-million-configuration Electro-Fenton study, even the most precise high-throughput system remains a “tabula rasa” that cannot leverage its own history. This results in a massive cold start penalty for every new campaign.

Our framework proves that learnability is the ultimate form

of sophistication. By internalising hand-feel and strategic patterns into long-term memory, our system transformed from an intelligent novice into an expert operator, completing a 30-day human-intensive campaign in just 45 hours. We maintain that the perceived unpredictability of a system is actually the manifestation of Empirical Expertise—the ability to navigate noise and stochasticity that rigid automation simply ignores. Learnability does not replace automation; it ensures that the automation matures through its own operations.

## 7. Challenges

Despite the transformative potential of the new paradigm for scientific discovery, realising the vision of self-evolving discovery imposes critical challenges that transcend the capabilities of current Fourth Paradigm platforms.

**Long-term Stability and Epistemic Corruption.** The first barrier is the tension between operational stability and error accumulation in high-throughput environments. Unlike the “tool metaphor” of the Fourth Paradigm, where systems are frequently reset, an evolving scientific agent must operate across extended temporal scales to achieve ontogenetic growth. In such scenarios, the inherent physical stochasticity of the laboratory, such as sensor drift, mechanical wear, and reagent degradation, can lead to a cumulative feedback loop. If the Embodied Learnability layer fails to adaptively calibrate against these drifting tolerances, the resulting evaluation noise ( $y$ ) propagates upward, polluting the system’s institutional memory. This leads to epistemic corruption, where the Cognitive Layer reinforces its intuition based on hardware artefacts rather than genuine chemical kinetics. Engineering a robust self-correction mechanism that maintains data integrity over lifelong operation remains a daunting task.

**Trustworthiness and the Black-Box Dilemma.** The second pivotal concern is the interpretability of model-driven discovery. As SDLs transition from rigid, instruction-following pipelines to open-ended systems governed by Bayesian learning and neural policies, a transparency gap emerges. Ensuring that the “strategic scent” evolved by the Experiment Layer and the reasoning refinement in the Cognitive Layer are scientifically grounded is essential for human, AI collaboration. The scientific community requires verifiable reasoning chains that bridge the gap between abstract vector representations and observable physical laws. Without this explainability, the novel configurations identified by the agent—particularly those deviating from conventional human intuition - may be dismissed as hallucinations or algorithmic artefacts rather than valid discoveries.

## 8. Conclusion

This position paper contends that the trajectory of autonomous scientific discovery is undergoing a fundamental phase transition, from the Fourth Paradigm of data-intensive science to the Fifth Paradigm of continuous learnability. We argue that this new paradigm requires learnability through embodied experimentation, not merely automation sophistication.

To operationalise this vision, we introduced the Trinity of Learnability, a unified framework that prevents the fragmentation of intelligence. By ensuring that Cognitive reasoning, Experimental optimisation, and embodied execution evolve in synchrony, our architecture overcomes the institutional amnesia that plagues current designs. The efficacy of this approach was substantiated by our Electro-Fenton case study, where the system bridged the reality gap to compress weeks of manual iteration into 45 hours of autonomous mastery. Ultimately, this work calls for a conceptual reimagining of the laboratory. We must cease viewing the SDL as a rigid, resettable tool and instead treat it as a learnable scientific organism capable of ontogenetic growth, which is a transformation essential for realising robust, end-to-end closed-loop autonomous discovery. In this emerging Fifth Paradigm, the value of an autonomous system is defined not by the volume of data it generates, but by the cumulative wisdom it inherits, marking the dawn of true AI-driven mastery in the physical sciences.

## References

- Adamo, A., Beingessner, R. L., Behnam, M., Chen, J., Jamison, T. F., Jensen, K. F., Monbaliu, J.-C. M., Myerson, A. S., Revalor, E. M., Snead, D. R., Stelzer, T., Weerapannant, N., Wong, S. Y., and Zhang, P. On-demand continuous-flow production of pharmaceuticals in a compact, reconfigurable system. *Science*, 352(6281):61–67, 2016. doi: 10.1126/science.aaf1337.
- Berridge, J. C. The application of the modified simplex method to the optimization of chromatographic conditions. *Journal of Chromatography A*, 244(1):1–14, 1982.
- Boiko, D. A., MacKnight, R., Kline, B., and Gomes, G. Autonomous chemical research with large language models. *Nature*, 624(7992):570–578, 2023.
- Box, G. E. P., Hunter, J. S., and Hunter, W. G. *Statistics for Experimenters: Design, Innovation, and Discovery*. Wiley Series in Probability and Statistics. Wiley-Interscience, Hoboken, N.J, 2nd ed edition, 2005. ISBN 978-0-471-71813-0.
- Bran, A. M., Cox, S., Schilter, O., Baldassari, C., White, A. D., and Schwaller, P. Augmenting large language models with chemistry tools. *Nature Machine Intelligence*, 6(5):525–535, 2024. doi: 10.1038/s42256-024-00832-8. URL <https://www.nature.com/articles/s42256-024-00832-8>.
- Burger, B., Maffettone, P. M., Gusev, V. V., Aitchison, C. M., Bai, Y., Wang, X., Li, X., Alston, B. M., Li, B., Clowes, R., et al. A mobile robotic chemist. *Nature*, 583(7815):237–241, 2020.
- Chen, S. and Jung, Y. Deep retrosynthetic reaction prediction using local reactivity and global attention. *JACS Au*, 1(10):1612–1620, 2021. doi: 10.1021/jacsau.1c00246. Published 2021, commonly cited as 2020 preprint.
- Christensen, M., Yunker, L. P. E., Adedeji, F., Häse, F., Roch, L. M., Gensch, T., dos Passos Gomes, G., Zepel, T., Sigman, M. S., Aspuru-Guzik, A., and Hein, J. E. Data-science driven autonomous process optimization. *Communications Chemistry*, 4(1):112, 2021. doi: 10.1038/s42004-021-00550-x. Autonomous optimization platform for chemical synthesis.
- Coley, C. W., Barzilay, R., Jaakkola, T. S., Green, W. H., and Jensen, K. F. Prediction of organic reaction outcomes using machine learning. *ACS Central Science*, 3(5):434–443, 2017. doi: 10.1021/acscentsci.7b00064.
- Coley, C. W., Thomas, D. A., Lummiss, J. A. M., Jaworski, J. P., Breen, C. P., Schultz, V., Hart, T. K., Fishman, J. S., Rogers, L., Gao, H., Hicklin, R. W., Plehiers, P. P., Byington, J., Piens, N., Struble, T. J., Chytil, M., Saikin, S. K., Hein, J. E., and Jensen, K. F. A robotic platform for flow synthesis of organic compounds informed by ai planning. *Science*, 365(6453):eaax1566, 2019. doi: 10.1126/science.aax1566. URL <https://doi.org/10.1126/science.aax1566>.
- Dai, T., Gongora, A. E., Bader, R., Peng, S., Roch, L. M., Aspuru-Guzik, A., and Cooper, A. I. Autonomous mobile robots for exploratory synthetic chemistry. *Nature*, 628(8008):590–596, 2024. doi: 10.1038/s41586-024-08173-7. URL <https://doi.org/10.1038/s41586-024-08173-7>.
- Darvish, K., Skreta, M., Zhao, Y., Yoshikawa, N., Som, S., Bogdanovic, M., Cao, Y., Hao, H., Xu, H., Aspuru-Guzik, A., Garg, A., and Shkurti, F. Organa: A robotic assistant for automated chemistry experimentation and characterization. *Matter*, 2025a. doi: 10.1016/j.matt.2024.10.017. URL [https://www.cell.com/matter/fulltext/S2590-2385\(24\)00542-3](https://www.cell.com/matter/fulltext/S2590-2385(24)00542-3). Advance online / In press (2024 preprint: arXiv:2401.06949).
- Darvish, K., Sohal, A., Mandal, A., Fakhrudeen, H., Radulov, N., Zhou, Z., Veeramani, S., Choi, J., Han, S., Zhang, B., et al. Matterix: toward a digital twin for

- robotics-assisted chemistry laboratory automation. *Nature computational science*, pp. 1–16, 2025b.
- Ding, A. W. and Li, S. Generative ai lacks the human creativity to achieve scientific discovery from scratch. *Scientific Reports*, 15(1):9587, 2025.
- Ding, K., Yu, J., Huang, J., Yang, Y., Zhang, Q., and Chen, H. Scitoolagent: a knowledge-graph-driven scientific agent for multitool integration. *Nature Computational Science*, pp. 1–11, 2025.
- Gao, F., Qian, X., Yu, X.-F., et al. A chemical autonomous robotic platform for end-to-end synthesis of nanoparticles. *Nature Communications*, 2025. doi: 10.1038/s41467-025-62994-2. URL <https://doi.org/10.1038/s41467-025-62994-2>.
- Garnett, R. *Bayesian Optimization*. Cambridge University Press, 2023.
- Gongora, A. E., Xu, B., Perry, W., Okoye, C., Riley, P., Reyes, K. G., Morgan, E. F., and Brown, K. A. A bayesian experimental autonomous researcher for mechanical design. *Science Advances*, 6(15):eaaz1708, 2020. doi: 10.1126/sciadv.aaz1708.
- Goodell, A. J., Chu, S. N., Rouholiman, D., and Chu, L. F. Large language model agents can use tools to perform clinical calculations. *npj Digital Medicine*, 8(1):163, 2025.
- Gordon, E. M., Barrett, R. W., Dower, W. J., Fodor, S. P. A., and Gallop, M. A. High-throughput synthesis and screening of combinatorial libraries. *Journal of Medicinal Chemistry*, 37(10):1385–1401, 1994.
- Granda, J. M., Donina, L., Dragone, V., Long, D.-L., and Cronin, L. Controlling an organic synthesis robot with machine learning to search for new reactivity. *Nature*, 559(7714):377–381, 2018. doi: 10.1038/s41586-018-0307-8.
- Gu, E., Tang, X., Langner, S., Duchstein, P., Zhao, Y., Levchuk, I., Hou, Y., Zhou, Y., Scheel, M. A., Osvet, A., Brabec, C. J., et al. Robot-based high-throughput screening of antisolvents for lead halide perovskites. *Joule*, 4(8):1806–1822, 2020. doi: 10.1016/j.joule.2020.06.013. URL <https://doi.org/10.1016/j.joule.2020.06.013>.
- Gupta, A., Savarese, S., Ganguli, S., and Fei-Fei, L. Embodied intelligence via learning and evolution. *Nature communications*, 12(1):5721, 2021.
- Hey, T., Tansley, S., Tolle, K. M., et al. *The fourth paradigm: data-intensive scientific discovery*, volume 1. Microsoft research Redmond, WA, 2009.
- Hie, B. L., Shanker, V. R., Xu, D., Bruun, T. U. J., Weidenbacher, P. A., Tang, S., Wu, W., Pak, J. E., and Kim, P. S. Efficient evolution of human antibodies from general protein language models. *Nature Biotechnology*, 42: 275–283, 2024. doi: 10.1038/s41587-023-01763-2.
- Huang, C., Chen, S., Li, Z., Qu, J., Xiao, Y., Liu, J., and Chen, Z. Geoagent: To empower llms using geospatial tools for address standardization. In *Findings of the Association for Computational Linguistics ACL 2024*, pp. 6048–6063, 2024.
- Häse, F., Roch, L. M., Kreisbeck, C., and Aspuru-Guzik, A. Phoenix: A bayesian optimizer for chemistry. *ACS Central Science*, 4(9):1134–1145, 2018. doi: 10.1021/acscentsci.8b00307.
- Isenhour, T. L. Robotics in the laboratory. *Journal of Chemical Information and Computer Sciences*, 25(3): 292–295, 1985.
- Jablonka, K. M., Schwaller, P., Ortega-Guerrero, A., and Smit, B. Leveraging large language models for predictive chemistry. *Nature Machine Intelligence*, 6:161–169, 2024. doi: 10.1038/s42256-023-00788-1.
- Jumper, J., Evans, R., Pritzel, A., et al. Highly accurate protein structure prediction with alphafold. *Nature*, 596(7873):583–589, 2021.
- Kang, Y., Seo, H., Lim, J., Kim, H. Q., Lee, S., Kang, M., Han, S. J., and Kim, J. Chatmof: an artificial intelligence system for predicting and generating metal-organic frameworks using large language models. *Nature Communications*, 15:4705, 2024. doi: 10.1038/s41467-024-49414-z.
- King, R. D., Whelan, K. E., Jones, F. M., Reiser, P. G. K., Bryant, C. H., Muggleton, S. H., Kell, D. B., and Oliver, S. G. Functional genomic hypothesis generation and experimentation by a robot scientist. *Nature*, 427(6971): 247–252, 2004. doi: 10.1038/nature02236. URL <https://doi.org/10.1038/nature02236>.
- King, R. D., Rowland, J., Oliver, S. G., Young, M., Aubrey, W., Byrne, E., Liakata, M., Markham, M., Pir, P., Soldatova, L. N., Sparkes, A., Whelan, K. E., and Clare, A. The automation of science. *Science*, 324(5923):85–89, 2009. doi: 10.1126/science.1165620. URL <https://doi.org/10.1126/science.1165620>.
- Krasner, J. Automation in the clinical laboratory. *Clinical Chemistry*, 13(7):560–574, 1967.
- Kreutz, J. E., Shukhaev, A., Du, W., Druskin, S., Daugulis, O., and Ismagilov, R. F. Evolution of catalysts directed by genetic algorithms in a plug-based microfluidic device tested with oxidation of methane by oxygen. *Journal of the American Chemical Society*, 132(9):3128–3132, 2010.

- Krishnadasan, S., Brown, R. J. C., deMello, A. J., and deMello, J. C. Intelligent routes to the controlled synthesis of nanoparticles. *Lab on a Chip*, 7(11):1434–1441, 2007.
- Lattimore, T. and Szepesvári, C. *Bandit Algorithms*. Cambridge University Press, 1 edition, July 2020. ISBN 978-1-108-57140-1 978-1-108-48682-8. doi: 10.1017/9781108571401.
- Li, B., Yan, T., Pan, Y., Luo, J., Ji, R., Ding, J., Xu, Z., Liu, S., Dong, H., Lin, Z., and Wang, Y. Mmedagent: Learning to use medical tools with multi-modal agent. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pp. 8745–8760, 2024.
- Li, J., Ding, C., Liu, D., Chen, L., and Jiang, J. Autonomous laboratories in china: an embodied intelligence-driven platform to accelerate chemical discovery. *Digital Discovery*, 2025.
- Lunt, A. M., Hou, X., Feng, J., et al. Modular, multi-robot integration of laboratories: an autonomous workflow for solid-state chemistry. *Chemical Science*, 15(7):2456–2463, 2024.
- Luo, X., Rechartd, A., Sun, G., Nejad, K. K., Yáñez, F., Yilmaz, B., Lee, K., Cohen, A. O., Borghesani, V., Pashkov, A., et al. Large language models surpass human experts in predicting neuroscience results. *Nature human behaviour*, 9(2):305–315, 2025.
- MacLeod, B. P., Parlane, F. G. L., Morrissey, T. D., Häse, F., Roch, L. M., Dettelbach, K. E., Moreira, R., Yunker, P. J., Aspuru-Guzik, A., Hein, J. E., and Berlinguette, C. P. Self-driving laboratory for accelerated discovery of thin-film materials. *Science Advances*, 6(20):eaaz8867, 2020. doi: 10.1126/sciadv.aaz8867. URL <https://doi.org/10.1126/sciadv.aaz8867>.
- Madani, A., Krause, B., Greene, E. R., Subramanian, S., Mohr, B. P., Holton, J. M., Olmos, J. L., Xiong, C., Sun, Z. Z., Socher, R., Fraser, J. S., and Naik, N. Large language models generate functional protein sequences across diverse families. *Nature Biotechnology*, 41:1099–1106, 2023. doi: 10.1038/s41587-022-01618-2.
- Maffettone, P. M., Friederich, P., Baird, S. G., Blaiszik, B., Brown, K. A., Campbell, S. I., Cohen, O. A., Davis, R. L., Foster, I. T., Haghmoradi, N., et al. What is missing in autonomous discovery: open challenges for the community. *Digital Discovery*, 2(6):1644–1659, 2023.
- Manzano, J. S., Kitson, P. J., Nguyen, D. T., Lai, Z., Baldoví, H. F., Gerard, C. J., and Cronin, L. An autonomous portable platform for universal chemical synthesis. *Nature Chemistry*, 14(11):1311–1318, 2022. doi: 10.1038/s41557-022-01016-w. URL <https://www.nature.com/articles/s41557-022-01016-w>.
- Matsuda, K., Ishii, T., and Ito, Y. Automated optimization of phase systems in high-speed countercurrent chromatography using a microcomputer. *Journal of Chromatography A*, 437:321–330, 1988.
- Mehr, S. H. M., Craven, M., Leonov, A. I., Keenan, G., and Cronin, L. A universal system for digitization and automatic execution of the chemical synthesis literature. *Science*, 370(6512):101–108, 2020. doi: 10.1126/science.abc2986.
- Merchant, A., Batzner, S., Schoenholz, S. S., Aykol, M., Cheon, G., and Cubuk, E. D. Scaling deep learning for materials discovery. *Nature*, 624:80–85, 2023. doi: 10.1038/s41586-023-06735-9.
- Merrifield, R. B. Automated synthesis of peptides. *Science*, 150(3693):178–185, 1965. doi: 10.1126/science.150.3693.178.
- Merrifield, R. B., Stewart, J. M., and Jernberg, N. Instrument for automated synthesis of peptides. *Analytical chemistry*, 38(13):1905–1914, 1966.
- Mittal, M., Yu, C., Yu, Q., Liu, J., Rudin, N., Hoeller, D., Yuan, J. L., Singh, R., Guo, Y., Mazhar, H., Mandekar, A., Babich, B., State, G., Hutter, M., and Garg, A. Orbit: A unified simulation framework for interactive robot learning environments. *IEEE Robotics and Automation Letters*, 8(6):3740–3747, 2023. doi: 10.1109/LRA.2023.3270034.
- Ni, B. and Buehler, M. J. Mechagents: Large language model multi-agent collaborations can solve mechanics problems, generate new data, and integrate knowledge. *Extreme Mechanics Letters*, 67:102131, 2024.
- Noé, F., Olsson, S., Köhler, J., and Wu, H. Boltzmann generators: Sampling equilibrium states of many-body systems with deep learning. *Science*, 365(6457):eaaw1147, 2019.
- Parrott, A. J., Bourne, R. A., Akien, G. R., Irvine, D. J., and Poliakoff, M. Self-optimizing continuous reactions in supercritical carbon dioxide. *Angewandte Chemie International Edition*, 49(39):7076–7080, 2010.
- Pyzer-Knapp, E. O., Pitera, J. W., Staar, P. W. J., et al. Accelerating materials discovery using artificial intelligence, high performance computing and robotics. *npj Computational Materials*, 8:84, 2022.
- Qu, Y., Huang, K., Yin, M., Zhan, K., Liu, D., Yin, D., Cousins, H. C., Johnson, W. A., Wang, X., Shah, M., et al. Crispr-gpt for agentic automation of gene-editing experiments. *Nature Biomedical Engineering*, pp. 1–14, 2025.

- Raccuglia, P., Elbert, K. C., Adler, P. D. F., Falk, C., Wenny, M. B., Mollo, A., Zeller, M., Friedler, S. A., Schrier, J., and Norquist, A. J. Machine-learning-assisted materials discovery using failed experiments. *Nature*, 533(7601): 73–76, 2016. doi: 10.1038/nature17439.
- Rainforth, T., Foster, A., Ivanova, D. R., and Smith, F. B. Modern Bayesian Experimental Design, November 2023.
- Rasmussen, C. E. and Williams, C. K. I. *Gaussian Processes for Machine Learning*. Adaptive Computation and Machine Learning. MIT Press, Cambridge, Mass., 3. print edition, 2008. ISBN 978-0-262-18253-9.
- Roch, L. M., Häse, F., Kreisbeck, C., Tamayo-Mendoza, T., Yunker, L. P. E., Hein, J. E., and Aspuru-Guzik, A. Chemos: Orchestrating autonomous experimentation. *Science Robotics*, 5(41):eaat5559, 2020. doi: 10.1126/scirobotics.aat5559.
- Schwaller, P., Vaucher, A. C., Laino, T., and Reymond, J.-L. Prediction of chemical reaction yields using deep learning. *Machine Learning: Science and Technology*, 2(1):015016, 2021. doi: 10.1088/2632-2153/abc81d.
- Segler, M. H. S., Preuss, M., and Waller, M. P. Planning chemical syntheses with deep neural networks and symbolic ai. *Nature*, 555(7698):604–610, 2018a. doi: 10.1038/nature25978.
- Segler, M. H. S., Preuss, M., and Waller, M. P. Planning chemical syntheses with deep neural networks and symbolic ai. *Nature*, 555(7698):604–610, 2018b. doi: 10.1038/nature25978. Same as segler2018, included for clarity in text citations.
- Seifrid, M., Hatrick-Simpers, J., Aspuru-Guzik, A., Kalinin, S. V., and Fung, V. Autonomous experimentation systems for materials development: A community perspective. *Matter*, 5(11):3657–3675, 2022. doi: 10.1016/j.matt.2022.09.012.
- Singhal, K., Azizi, S., Tu, T., Mahdavi, S. S., Wei, J., Chung, H. W., Scales, N., Tanwani, A., Cole-Lewis, H., Pfohl, S., Payne, P., Seneviratne, M., Gamble, P., Kelly, C., Babiker, A., Schärli, N., Chowdhery, A., Mansfield, P., Demner-Fushman, D., Agüera y Arcas, B., Webster, D., Corrado, G. S., Matias, Y., Chou, K., Gottweis, J., Tomasev, N., Liu, Y., Rajkomar, A., Barral, J., Semturs, C., Karthikesalingam, A., and Natarajan, V. Large language models encode clinical knowledge. *Nature*, 620:172–180, 2023. doi: 10.1038/s41586-023-06291-2.
- Smallwood, H. Design of experiments in industrial research. *Analytical Chemistry*, 19(12):950–952, 1947.
- Smith, A. A., Wong, E. L., Donovan, R. C., Chapman, B. A., Harry, R., Tirandazi, P., Kanigowska, P., Gendreau, E. A., Dahl, R. H., Jastrzebski, M., et al. Using a gpt-5-driven autonomous lab to optimize the cost and titer of cell-free protein synthesis. *bioRxiv*, pp. 2026–02, 2026.
- Srinivas, N., Krause, A., Kakade, S. M., and Seeger, M. Gaussian Process Optimization in the Bandit Setting: No Regret and Experimental Design. *IEEE Transactions on Information Theory*, 58(5):3250–3265, May 2012. ISSN 0018-9448, 1557-9654. doi: 10.1109/TIT.2011.2182033.
- Steiner, S., Wolf, J., Glatzel, S., Andreou, A., Granda, J. M., Keenan, G., Hinkley, T., Aragon-Camarasa, G., Kitson, P. J., Angelone, D., et al. Organic synthesis in a modular robotic system driven by a chemical programming language. *Science*, 363(6423):eaav2211, 2019.
- Szymanski, N. J., Zitting, L., Owens, S. G., et al. An autonomous laboratory for the accelerated synthesis of novel materials. *Nature*, 624(7975):86–91, 2023.
- Thirunavukarasu, A. J., Ting, D. S. J., Elangovan, K., Gutierrez, L., Tan, T. F., and Ting, D. S. W. Large language models in medicine. *Nature Medicine*, 29:1930–1940, 2023. doi: 10.1038/s41591-023-02448-8.
- Todorov, E., Erez, T., and Tassa, Y. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 5026–5033, 2012. doi: 10.1109/IROS.2012.6386109.
- Topham, S. A. Automated catalyst testing in the chemical industry. *Chemistry and Industry*, pp. 705–710, 1985.
- Wang, J. Memento 2: Learning by Stateful Reflective Memory, December 2025.
- Williams, K. S., King, R. D., Oliver, S. G., et al. Cheaper, faster drug development validated by the repositioning of drugs against neglected tropical diseases. *Journal of The Royal Society Interface*, 12(104):20141289, 2015. doi: 10.1098/rsif.2014.1289. URL <https://doi.org/10.1098/rsif.2014.1289>.
- Wu, H., He, Z., Zhang, X., Yao, X., Zheng, S., Zheng, H., and Yu, B. Chateda: A large language model powered autonomous agent for eda. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 43(10):3184–3197, 2024.
- Xie, Y., Jiang, B., Mallick, T., Bergerson, J., Hutchison, J. K., Verner, D. R., Branham, J., Alexander, M. R., Ross, R. B., Feng, Y., et al. Marsha: multi-agent rag system for hazard adaptation. *npj Climate Action*, 4(1):70, 2025.
- Yang, Z., Zeng, X., Zhao, Y., and Chen, R. AlphaFold2 and its applications in the fields of biology and medicine. *Signal Transduction and Targeted Therapy*, 8(1):115, 2023.

Zhang, Z., Ren, Z., Hsu, C.-W., Chen, W., Hong, Z.-W., Lee, C.-F., Penn, A., Xu, H., Zheng, D. J., Miao, S., et al. A multimodal robotic platform for multi-element electrocatalyst discovery. *Nature*, 647(8089):390–396, 2025.

Zhao, H., Chen, W., Huang, H., Sun, Z., Chen, Z., Wu, L., Zhang, B., Lai, F., Wang, Z., Adam, M. L., Pang, C. H., Chu, P. K., Lu, Y., Wu, T., Jiang, J., Yin, Z., and Yu, X.-F. A robotic platform for the synthesis of colloidal nanocrystals. *Nature Synthesis*, 2(6):505–514, 2023a. doi: 10.1038/s44160-023-00250-5. URL <https://doi.org/10.1038/s44160-023-00250-5>.

Zhao, T. Z., Kumar, V., Levine, S., and Finn, C. Learning fine-grained bimanual manipulation with low-cost hardware, 2023b. URL <https://arxiv.org/abs/2304.13705>.

Zhao, X., Li, T., Wu, C., et al. Biomars: Bimanual llm-driven robotic scientist for autonomous cell culture, 2025. URL <https://arxiv.org/abs/2507.01485>.

Zhou, H., Chen, Y., Guo, S., Yan, X., Lee, K. H., Wang, Z., Lee, K. Y., Zhang, G., Shao, K., Yang, L., et al. Memento: Fine-tuning llm agents without fine-tuning llms. *arXiv preprint arXiv:2508.16153*, 2025.

Zhu, Q., Huang, Y., Xu, Z., Chen, Z., Zhang, Y., Zhao, H., Li, J., Wang, B., Zeng, X., Luo, Y., Jiang, J., and Shang, W. Automated synthesis of oxygen-producing catalysts from martian meteorites by a robotic ai chemist. *Nature Synthesis*, 3(4):319–329, 2024. doi: 10.1038/s44160-023-00424-1. URL <https://www.nature.com/articles/s44160-023-00424-1>.

Zymark Corporation. Zymate laboratory automation system, 1982. First presented at Pittsburgh Conference, 1982.

## A. Hardware Implementation of the SDL Prototype

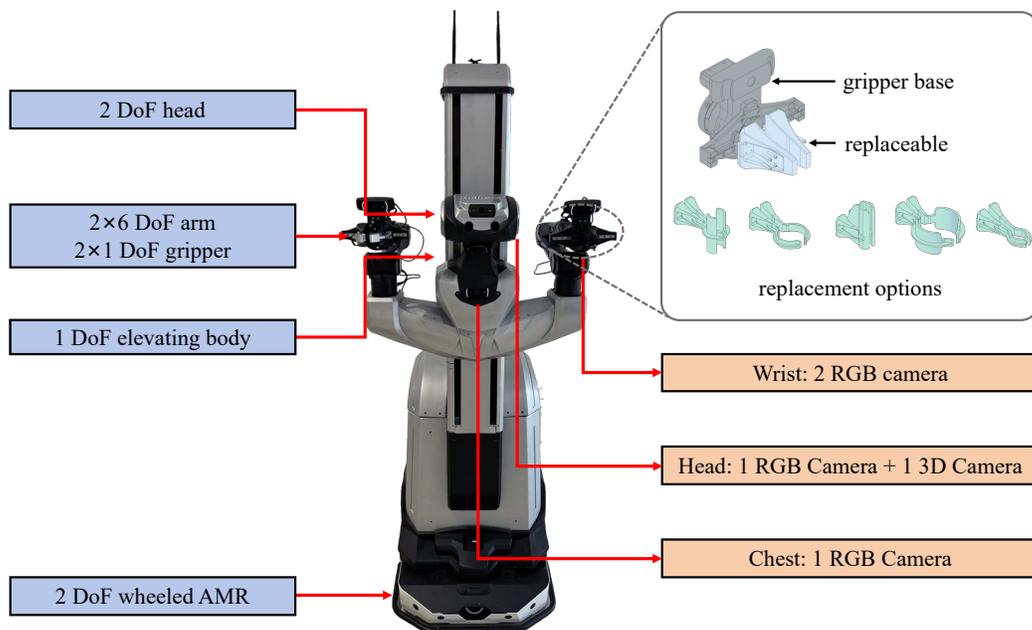


Figure 7. Overview of the robot platform of the prototype. This dual-arm robot comprises two 6-DoF robotic arms equipped with 1-DoF grippers, a 2-DoF articulated head, an omnidirectional AMR base offering 2 planar DoFs for holonomic motion, and a 1-DoF lift for vertical adjustment. These degrees of freedom are highlighted in the blue box. The platform also integrates a synchronised multi-camera sensing suite consisting of one RGB and one depth camera mounted on the head, two RGB cameras on each wrist, and one RGB camera on the chest, as indicated in the orange box. Together, these components provide rich proprioceptive and multi-view visual feedback to support both global perception and fine-grained manipulation.

To validate the proposed framework and the concept of Tripartite Learnability, we constructed a physical prototype comprising a mobile manipulator agent and a suite of laboratory instrumentation.

### A.1. Mobile Manipulation Platform

The agent is a dual-arm mobile manipulator designed for holonomic navigation and dexterous manipulation within a standard laboratory setting, as illustrated in Figure 7. The mobility is provided by a 2-DoF omnidirectional Autonomous Mobile Robot (AMR) base, coupled with a 1-DoF vertical lift torso to expand the vertical workspace. The manipulation system consists of two 6-DoF robotic arms, mimicking human bimanual operation. Each arm is terminated with a 1-DoF parallel gripper featuring replaceable fingers to adapt to various glassware geometries.

To support the *Embodied Learnability* layer, the robot is equipped with a comprehensive multi-view sensing suite. This includes a 2-DoF articulated head containing a high-resolution RGB camera and a 3D depth camera for global perception. For fine-grained manipulation and visual servoing, additional RGB cameras are mounted on the chest and both wrists. This configuration provides occlusion-robust visual feedback necessary for complex tasks such as pipetting and pouring.

For control and policy learning, the dual arm robot’s dexterous skills are powered by Action Chunking with Transformers (ACT) framework (Zhao et al., 2023b), trained on a dataset of 480 human-demonstrated trajectories featuring randomised object poses and visual distractors to ensure robust Embodied Learnability. These policies are organised into a Visuomotor Policy Dictionary (VPD), enabling the high-level agent to invoke specific manoeuvres via semantic retrieval and monitor task completion through a real-time joint-angle similarity-based feedback mechanism.

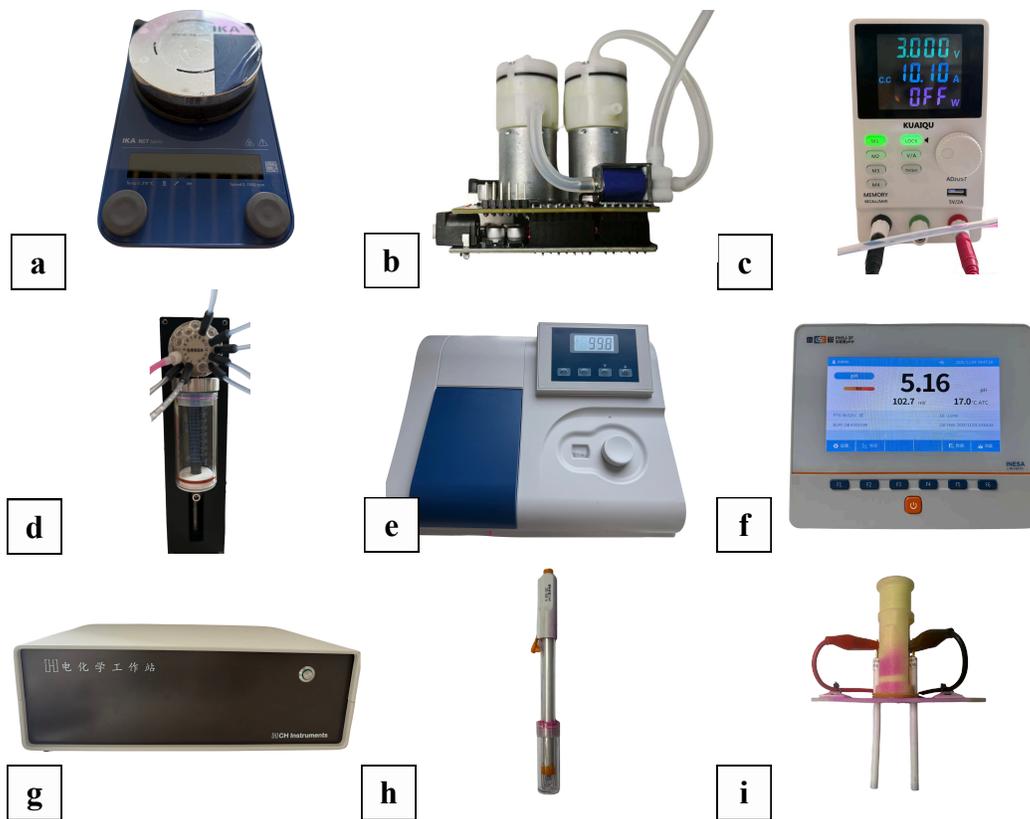


Figure 8. Experimental devices used in our laboratory: (a) IKA RCT Basic Magnetic Stirrer, (b) Arduino Uno R3 Air Pump, (c) QUAQU SPPS3010D-232 Programmable Power Supply, (d) ZHIGAO DP-03-60 Syringe Pump, (e) INESA 721G Spectrophotometer, (f) LEICI PHSJ-3F pH Meter, (g) Chenhua CHI660F Electrochemical Workstation, (h) Probe of the LEICI PHSJ-3F pH Meter, (i) Custom Electrode Assembly.

## A.2. Experimental Instrumentation

The laboratory environment is populated with a diverse array of analytical and executive devices, shown in Figure 8. The hardware selection targets general chemistry workflows, specifically solution preparation and characterisation. The hardware ensemble is categorised into executive and analytical units to support a complete, end-to-end discovery cycle:

- **Executive Hardware:** For synthesis and sample processing, the setup includes an IKA RCT Basic magnetic stirrer (a) for mixing, a QUAQU SPPS3010D programmable power supply (c) for electrochemical driving, and fluid handling systems comprising an Arduino-controlled air pump (b) and a ZHIGAO syringe pump (d).
- **Analytical Hardware:** Characterisation is performed using an INESA 721G spectrophotometer (e) for optical density measurements, a LEICI PHSJ-3F pH meter (f) with a corresponding probe (h), and a Chenhua CHI660F electrochemical workstation (g).
- **Custom Interface:** To bridge the gap between the robot and standard equipment, we designed custom fixtures, such as the electrode assembly (i), which allows the robot to handle delicate sensors robustly.

## B. The GP-UCB Implementation of BO

In this scenario, we implement GP-UCB (Srinivas et al., 2012), a classic version of BO based on GPs (Rasmussen & Williams, 2008) and the upper-confidence-bound (UCB) (Lattimore & Szepesvári, 2020) acquisition function.

GPs are the most widely used surrogate models in BO due to their flexibility in regressing arbitrary unknown functions with small datasets. Conceptually, a GP is a collection of random variables (indexed by  $\mathcal{X}$ , in our case), where any  $n$  variables follow an  $n$ -variate Gaussian distribution. To specify a GP, we need a mean function  $m : \mathcal{X} \mapsto \mathbb{R}$  and a covariance kernel

$\kappa : \mathcal{X}^2 \mapsto \mathbb{R}$ , denote as  $\mathcal{GP}(m, \kappa)$ :

$$\mu(x) = \mathbb{E}f(x), \quad \kappa(x, x') = \text{Cov}(f(x), f(x')). \quad (1)$$

**GP Prior and Initial Data.** The prior  $P(f)$  is defined by selecting a prior mean  $m(x)$  and prior kernel  $k(x, x')$ . A common choice is  $m(x) = 0$ . The prior kernel  $k$  is chosen based on the properties of the decision space  $\mathcal{X}$  and the available knowledge about the objective function  $f$ . For example, for a smooth  $f$  on a closed set in Euclidean space  $\mathbb{R}^d$ , a plausible choice is the radial basis function (RBF) kernel:

$$k(x, x') = c \cdot \exp \left\{ -\|x - x'\|^2 / 2\ell^2 \right\}, \quad (2)$$

where  $c$  and  $\ell$  are hyperparameters. To acquire the initial dataset  $\mathcal{D}_0$ , we uniformly draw  $m$  random parameters from  $\mathcal{X}$  and evaluate their objectives.

**GP Posterior Surrogate.** Assume that each observation  $y_t = f(x_t) + \epsilon_t$  is perturbed by  $\epsilon_t \sim N(0, \sigma^2)$  i.i.d. Gaussian noise. Given that  $P(f) \sim \mathcal{GP}(m(x), k(x, x'))$  is the prior distribution of  $f$ , it can be derived that the posterior distribution of  $f$  given  $\mathcal{D}_t = \{(x_i, y_i)\}_{i=1}^t$  is also a GP, with mean function  $\mu_t$  and kernel  $k_t$ . For simplicity, we define  $\mathbf{y}_t := (y_1, \dots, y_t)^\top$ ,  $\mathbf{k}_t(x) := (k(x, x_1), \dots, k(x, x_t))^\top$ ,  $\mathbf{m}_t := (m(x_1), \dots, m(x_t))^\top$ , and  $\mathbf{K}_t := (k(x_i, x_j))_{i,j}^{t \times t}$ . The posterior distribution is given by (Garnett, 2023):

$$\mu_t(x) = \mathbf{k}_t(x)^\top (\mathbf{K}_t + \sigma^2 \mathbf{I})^{-1} (\mathbf{y}_t - \mathbf{m}_t), \quad (3)$$

$$\kappa_t(x, x') = k(x, x') - \mathbf{k}_t(x)^\top (\mathbf{K}_t + \sigma^2 \mathbf{I})^{-1} \mathbf{k}_t(x'). \quad (4)$$

Consequently, we have the surrogate  $\hat{f}_t \sim \mathcal{GP}(\mu_t, \kappa_t)$ . Note that in our current implementation, we only have one performance criterion, which is exactly the result of interest:

$$\ell(y) = -y \in \mathbb{R}, \quad f_\pi(x) = -\mathbb{E}_{y \sim \nu, \gamma \sim \pi(\cdot|x)}[\ell(y)] = y.$$

Therefore, this GP surrogate for  $f_\pi$  is also the agent's modelling  $\hat{h}$  as introduced in Section ??.

**Acquisition Functions.** The GP-UCB algorithm defines the AF as the upper confidence bound:

$$\alpha(x|\mathcal{D}_t) = \mu_t(x) + \beta \sigma_t(x), \quad (5)$$

where  $\sigma_t(x) = \sqrt{\kappa_t(x, x)}$  is the posterior standard deviation, and  $\beta$  is a hyperparameter. By setting a positive  $\beta$ , the UCB encourages exploration where the posterior surrogate is uncertain, while trading off the exploitation of high expected values.

**AF Optimizer.** To balance computational efficiency and precision, we adopt a hierarchical grid search on the UCB acquisition function. It starts with a coarse grid of low resolution on  $\mathcal{X}$ , and recursively searches the regions near the high-AF points using a denser grid. The maximum of the AF is returned once a designated precision is reached. We also note that the decision space  $\mathcal{X}$  must be discretised based on the manipulatory accuracy of hardware. Therefore, the final output of the AF optimiser is mapped to the closest point in the discretised  $\mathcal{X}$ .