
Task Alignment Outweighs Framework Choice in Scientific LLM Agents

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 Large language models (LLMs) equipped with external tools through agentic frame-
2 works promise to overcome domain-specific limitations by providing specialized
3 capabilities for scientific applications. However, the extent to which these systems
4 genuinely enhance performance in complex scientific domains remains poorly
5 understood. Here we present *Corral*, a modular benchmarking framework that sys-
6 tematically evaluates LLM-based agents across four expert-designed environments
7 spanning molecular dynamics, machine learning, catalysis, and spectroscopy in
8 chemistry and materials science. Through comprehensive evaluation of state-of-the-
9 art models using different agent scaffolds, we demonstrate that the choice of agentic
10 framework—whether ReAct or tool-calling—plays a surprisingly minor role in
11 determining success. Instead, performance depends critically on the semantic align-
12 ment between available tools and task requirements, measured through embedding
13 similarity. When this alignment is poor, even sophisticated reasoning frameworks
14 cannot compensate for inadequate tool provisioning and a lack of domain knowl-
15 edge. Conversely, when base models possess sufficient domain knowledge, agentic
16 frameworks can introduce unnecessary overhead without meaningful benefits. Our
17 findings challenge the assumption that agentic systems provide a universal solution
18 to model limitations, revealing instead that, currently, successful scientific agents
19 might require the same level of domain expertise in tool design that was promised
20 to be circumvented.

21 1 Introduction

22 Large language models (LLMs) excel at generating text but struggle when tasks require interacting
23 with external systems, accessing real-time information, or executing multi-step procedures [1, 2].
24 Agent frameworks have been proposed to address these limitations by enabling models to plan
25 sequences of actions, use specialized tools, and operate within specific environments [3, 4, 5, 6,
26 7, 8, 9, 10]. This approach might be particularly critical in scientific applications, where precise
27 domain knowledge is essential for accurate decision-making [11, 12], in which base models show
28 limitations [13, 14]. However, the extent to which these tools genuinely enhance model performance
29 remains unclear. While initiatives like the Model Context Protocol (MCP) [15] have improved
30 tool accessibility, the mere provision of additional tools can be counterproductive; tools may act as
31 distractors rather than valuable assets, leading to inefficient or erroneous outputs [16, 17].

32 Current evaluation methodologies predominantly focus on high-level metrics such as overall per-
33 formance, tool usage count, and execution time, with an emphasis on coding and machine learning
34 tasks [18]. Although essential for assessing production-ready systems, these metrics often overlook
35 some of the critical components of the systems and do not provide practical guidance on how to
36 improve them [19]. By focusing solely on the end result, without considering the process of tool uti-

lization, these evaluations limit our understanding of the models’ true capabilities and can even result in agents failing to fully comprehend the functions available to them [20]. This lack of granularity is particularly problematic in scientific fields like chemistry or biology, where agents utilize tools not only for retrieval, but also for experiment execution, and knowledge synthesis [21].

Moreover, the importance of task-tool alignment is not properly understood. It is often unclear whether the provided tools are truly relevant or if the agent understands when and how to use them. Here, we demonstrate that the alignment between tasks and tools is indicative of agentic performance. To do so, we built a novel modular benchmarking framework, systematically implemented multiple scientific environments and agents, and conducted ablation studies.




Figure 1: **Representation of the core contribution in this work.** We present the *Corral* framework, which hosts environments on a server separating them from the agents, enabling an evaluation of how the different components affect the results. This analysis of the different components allows us to make actionable recommendations on how to improve LLM-based agents.

As illustrated in Figure 1, our main contributions are:

- We present *Corral*, a modular framework to implement, run, and test LLM-based agents. Within *Corral*, four expert-designed environments are used to assess the performance of current models.
- We benchmark two SOTA LLMs and two of the most widely implemented agent scaffolds, providing a comprehensive evaluation of current agent capabilities.
- We perform an analysis of the factors that determine the success of an LLM-based agent. In particular, we propose the use of embedding similarity between tools and tasks to estimate agents’ performance. The alignment between tools and tasks is especially important when the agent lacks the foundational knowledge required to carry out the task.

A detailed section on related research is presented in Appendix A.

2 *Corral* : Agent Evaluation Framework

The *Corral* framework, represented in Figure C.1, is designed to define, execute, and evaluate agent capabilities within different isolated environments (code snippets demonstrating how to implement the different components in the framework are shown in Appendix E). Inspired by *Aviary* [22] and prior works [23, 24, 25], we adopt a model similar to a Language Decision Process (LDP), which is a specific formulation of a Partially Observable Markov Decision Process (POMDP) where observations and actions are represented in text.

64 2.1 The Environment

65 The Environment ($\mathcal{E}$) is responsible for serving the task guide to the agent (which includes the task
66 description and input data from previous tasks, if any), providing the necessary tools, and evaluating
67 the agent’s performance. Formally, the environment implements the core components of the LDP.

68 **State Space ($\mathcal{S}$).** The state s in $\mathcal{S}$ represents a complete snapshot of the environment at a given
69 time step. In *Corral*, this is captured by the `TaskState` object. A state s_t at time t is a tuple:

$$s_t = (P, M_t, C_t, d, \sigma) \in \mathcal{S} \quad (1)$$

70 where:

- 71 • P is the static task prompt or problem description.
- 72 • $M_t = \{m_1, m_2, \dots, m_t\}$ is the set of messages exchanged between the agent and the
73 environment.
- 74 • $C_t = \{c_1, c_2, \dots, c_k\}$ is the history of k tool calls made up to time t . Each c_i is a `ToolCall`
75 object containing the tool name, arguments, result, and status.
- 76 • d is a boolean flag indicating if the task is completed (`is_completed`).
- 77 • σ is the final score, which is defined only when d is true.

78 During benchmarking, d and σ are not observable to the agent. However, this could be used to
79 optimize the system interacting with the environment during training.

80 **Tools.** A key responsibility of the environment is provisioning a set of tools $\mathcal{T}_{env}$. Each tool
81 $\tau \in \mathcal{T}_{env}$ is a function that maps a set of arguments to a result returned as a string.

82 **Task ($\mathcal{K}$).** A task $\mathcal{K}$ defines a specific, self-contained problem. It is formally a tuple:

$$\mathcal{K} = (P, \mathcal{T}_{req}, f_{score}, I_0) \quad (2)$$

83 where:

- 84 • P is the task prompt given to the agent.
- 85 • $\mathcal{T}_{req} \subseteq \mathcal{T}_{env}$ is the specific subset of tools required for this task.
- 86 • f_{score} is the custom scoring function used to implement the reward function R .
- 87 • I_0 is the initial input data for the task.

88 **Task Group ($\mathcal{G}$).** A task group allows for the creation of complex workflows by defining dependen-
89 cies between tasks. Formally, a task group is a directed graph:

$$\mathcal{G} = (V, E) \quad (3)$$

90 where $V = \{\mathcal{K}_i\}_{i=1}^N$ is a set of N tasks serving as the vertices of the graph. The set $E \subseteq V \times V$
91 contains the directed edges, where an edge $(\mathcal{K}_i, \mathcal{K}_j) \in E$ implies that the output of task $\mathcal{K}_i$ is a
92 required input for task $\mathcal{K}_j$. The directed nature of the graph guarantees a well-defined execution
93 order.

94 2.2 The Agent

95 The Agent is the decision-making entity that interacts with the environment to solve a task. It is
96 implemented as a policy π that maps the history of observations to a probability distribution over the
97 available actions.

98 **Action Space ($\mathcal{A}$).** The action space $\mathcal{A}$ consists of all possible actions the agent can take. This
99 space is the union of all possible tool calls and a special submission action, `submit`.

$$\mathcal{A} = \left(\bigcup_{\tau_i \in \mathcal{T}_{env}} \text{call}(\tau_i, \text{args}_i) \right) \cup \text{submit}(\text{answer}) \quad (4)$$

100 where args_i represents a set of arguments for tool τ_i . An action $a_t \in \mathcal{A}$ is thus either a tool invocation
101 or the submission of a final answer.

102 **Policy (π).** The policy π selects an action a_t at time t based on the history of observations $h_t =$
 103 $(o_1, o_2, \dots, o_t)$.

$$a_t \sim \pi(\cdot | h_t) \quad (5)$$

104 The $\cdot$ stands for any possible action a from the entire action space $\mathcal{A}$. The *Corral* framework is
 105 agnostic to the internal reasoning of the agent, allowing for different implementations of π , such as
 106 ReAct [26], planning algorithms [27], or simple tool-calling models.

107 2.3 Interaction Dynamics

108 The *Corral* runner orchestrates the agent-environment interaction over a series of tasks (Figure C.1
 109 illustrates the *Corral* framework). Each task is part of a separate and isolated environment on the
 110 *Corral* server. Within its corresponding isolated environment, the agent’s interaction unfolds over
 111 a sequence of discrete time steps, concluding only when it executes the `submit` action. Upon this
 112 terminal action, the *Corral* runner finalizes the interaction with the current environment and proceeds
 113 to the next task.

114 **Transition Function (T).** The state transition function $T : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{S}$ describes how the state
 115 evolves.

$$s_{t+1} = T(s_t, a_t) \quad (6)$$

116 If $a_t = \text{call}(\tau, \text{args})$, the environment executes the tool and updates the state by appending a new
 117 `ToolCall` record to the history C_t , resulting in C_{t+1} . If $a_t = \text{submit}(\text{answer})$, the environment
 118 updates the state by setting d to true and computing a score.

119 **Observation Space ($\mathcal{O}$) and Function (Z).** After the environment transitions to a new state s_{t+1} ,
 120 the agent receives an observation $o_{t+1} \in \mathcal{O}$. The observation function $Z : \mathcal{S} \rightarrow \mathcal{O}$ could be stochastic
 121 depending up on the tool.

$$o_{t+1} = Z(s_{t+1}) \quad (7)$$

122 The observation is typically the string-based result or error message from the most recent tool call,
 123 which corresponds to the `result` or `error_message` field in the last element of the tool call history
 124 C_{t+1} .

125 **Reward Function (R).** The reward function $R : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ provides the feedback signal to the
 126 agent. In *Corral*, during benchmarking, rewards are sparse and terminal. A reward is only issued
 127 when the agent takes the `submit` action.

$$R(s_t, a_t) = \begin{cases} f_{\text{score}}(\text{answer}) & \text{if } a_t = \text{submit}(\text{answer}) \\ 0 & \text{otherwise} \end{cases} \quad (8)$$

128 The reward is the agent’s score in that particular task. In our experiments, we stick to binary rewards
 129 (either 0 or 1).

130 3 Results

131 The environments for our evaluations, detailed in Appendix F, are deliberately designed to represent
 132 a spectrum of complexity and require specialized knowledge. The ML and `OPENCATALYST` tasks
 133 are more aligned with standard computational workflows, while molecular dynamics (MD) requires
 134 nuanced scientific knowledge, and especially `SPECTRA` requires reasoning over and validating the
 135 hypotheses it generates based on chemical observations.

136 To disentangle the effect of the framework in the policy, we compare two agent strategies, ReAct [26]
 137 and tool-calling. ReAct promotes an iterative *thought-action-observation* loop. The tool-calling
 138 agent is a more direct, function-execution approach, which might be more efficient for well-defined
 139 workflows. We probe the capability of the core reasoning engine of the policy (π) by testing two
 140 different foundation models.

3.1 Overall Performance

Figure 2 shows the agents’ ability to successfully complete the task within five attempts (pass@5, metrics used in the work are detailed in Appendix B). The illustration provides a high-level view of their capabilities across the four scientific environments, which are ranked by the domain expertise required to solve the tasks. The SPECTRA environment requires most domain-specific skills, the ML environment the least. The results underscore the significant impact of domain complexity on the policy’s (π) ability to achieve a high reward, R . Interestingly, these results follow an inverse trend compared to the number of tool calls, as reported in Appendix H, suggesting that agents tend to sample more tools when unsure of how to solve the tasks.




Figure 2: **Impact of domain expertise on agent performance across environments.** This plot shows the pass@5 scores for each agent across the four environments, ordered by the level of required domain expertise (from low at the bottom to high at the top). There is a clear inverse correlation: as the need for specialized knowledge increases, the agents’ ability to successfully complete the task within five attempts decreases. This visualizes the significant challenge that deep domain expertise poses to the agents’ policies.

Performance and Domain Familiarity Comparison In the OPENCATALYST and ML environments, both Claude-3.5 Sonnet and GPT-4o achieve near-perfect scores. This suggests that for these tasks, the policies (π) can effectively map the initial task prompt P to the correct sequence of tool calls. Conversely, environments demanding deep, specialized knowledge as MD and SPECTRA, result in lower scores. The particularly low scores in the SPECTRA environment highlight the difficulty the policy faces in interpreting chemical data to select the optimal sequence of actions from $\mathcal{A}$.

Agent and Model Comparison In the more demanding MD and SPECTRA environments, the differences between agents are very small. This is surprising since the tool-calling agent is a much simpler implementation that does not enforce agents’ reasoning compared with ReAct. However, in other environments where we suspect that the policies are trained on the reasoning chain in ReAct, it turns out to be helpful. Claude-3.5 Sonnet consistently shows a slight performance advantage over GPT-4o, especially in these challenging domains. This might suggest that Claude-3.5 Sonnet is more familiar with chemistry and material science compared to GPT-4o, consistent with prior research [14, 13].

3.2 Decomposing Tasks

To validate our hypothesis that the policy, π , is less likely to converge on a successful action trajectory for unfamiliar tasks, we evaluate the tasks in two different formats. As *single* tasks ($\mathcal{K}$), and *chained*




Figure 3: **Performance comparison between single and chained tasks.** This plot shows the pass@5 scores when the task is decomposed into simple one-step tasks, and when it is not. All tools are in COMPREHENSIVE verbosity (refer to Appendix D for a detailed explanation about the different levels of tool descriptions considered). In the more challenging MD and SPECTRA environments, all agent configurations show a significant increase in overall score when the task is presented as a chained sequence of sub-tasks. Table comparing results for all different tool verbosity in Table J.11

tasks, where the *single* tasks are broken down to smaller single-step problems, forming a task group $\mathcal{G}$. In a chained configuration, dependencies are defined and the output of task $\mathcal{K}_i$ serves as the initial input, I_0 , for task $\mathcal{K}_j$ (see example in Appendix F.4.3). By breaking down the tasks into smaller tasks, the action space (or the action trajectory) over which the policies have to reason is constrained.

Improved Performance through Decomposition in Complex Domains As shown in Figure 3, for the MD and SPECTRA environments, all agent and model configurations show a significant improvement in performance when moving from a *single* task to a *chained* one. For example, in the SPECTRA environment, the Claude-3.5 Sonnet tool-calling agent’s score jumps from below 0.25 on the single task to nearly 0.50 on the chained task. This strongly suggests that decomposing a complex problem into a structured sequence of smaller tasks provides a crucial scaffold that helps agents succeed if many domain skills are needed.

Performance Breakdown by Task Categories To further ground our understanding that lower performance could be attributed to the inability of the model to reason and validate scientific hypotheses and tasks, we classified all the decomposed tasks into a set of categories (“retrieval”, “code execution”, “experiment execution”, “reasoning”, and “validation”). For example, fetching a structure from *Materials Project* (MP)[28] in the OPENCATALYST environment is classified as “retrieval”, while running a Large-scale Atomic/Molecular Massively Parallel Simulator (LAMMPS) simulation is classified as “experiment execution”. This classification was done manually by the experts who designed the environments.

Figure 4 plots the average pass@5 score for each of the categories. This breakdown reveals that the primary challenge for current agents lies in higher-order reasoning, validation of the generated history, and execution rather than simple tool use. Again, as for the single task results, we observe that both agents, ReAct and tool-calling, perform very similarly despite the difference in complexity.

3.3 Tool Verbosity Ablation

It is well known that the amount of context provided to an LLM is important for the performance [29, 30]. In an attempt to systematically understand the importance of this for the agent performance, we encoded the tool descriptions in a structured format (described in Appendix D) and ablated the effect of only showing parts of this to the model. We evaluated agent performance across three progressively detailed verbosity levels: BRIEF, WORKFLOW, and COMPREHENSIVE. This tiered approach allows us



Figure 4: **Performance breakdown by task categories.** This figure illustrates the average pass@5 scores of different agent configurations across five task categories. All tasks were manually classified into different categories by the experts who designed the environments. While agents excel at concrete tasks like “retrieval” and “code execution”, performance declines for more complex categories such as “experiment execution”, “reasoning”, and “validation”. The average score here is the mean of pass@5 scores for each of the chained tasks.




Figure 5: **Results comparing single, chained, and the ablation for the Spectra environment.** The figure shows that by incorporating a new tool, the distance between tools and tasks is closed, resulting in a massive improvement in the scores of the single task being comparable to the ones in the chained task.

206 to measure how performance scales as the agent is supplied with increasingly rich contextual and
 207 operational knowledge about its tools.

208 **Level of Description Is Not All You Need** As detailed in Table J.11, for `OPENCATALYST` and `ML`
 209 environments, performance is consistently high for the three levels of tool description that were
 210 considered. For the most complex environments, `SPECTRA` and `MD`, the scores, despite being worse
 211 than for the other environments, do not show a clear trend in which verbosity is better. For example,
 212 in the `SPECTRA` environment, it is possible to see that the best score overall is achieved in the `BRIEF`
 213 description scenario by the tool-calling agent in the chained fashion. However, that same agent
 214 has a very similar performance when the tool descriptions follow the `WORKFLOW` level. Similarly,
 215 when evaluating the full or single task, the best score is obtained with the `COMPREHENSIVE` level of
 216 description.

207 3.4 Dependence on Task and Tool Similarity

208 In our analysis, we observe that the `SPECTRA` environment performs much worse than the rest of
 209 the environments, independent of the level of detail in the tool description. We hypothesize that
 210 this is because the tools provided in this environment are less aligned with the most difficult step
 211 in the tasks compared to other environments. To probe this hypothesis, we embedded the tools
 212 and task descriptions and measured their distances. A detailed description of how the embeddings
 213 are obtained and the distances calculated is provided in Appendix K. Table 1 shows that `SPECTRA`
 214 shows the largest mean minimum distances between tasks and tools, followed by the `MD`, `ML`, and
 215 `OPENCATALYST` environments. This ordering correlates with the overall performance of the agents
 216 in these environments. Additionally, `SPECTRA` presents the second-highest mean distance between
 217 tasks, only after `ML`. We hypothesize that the larger distances between consecutive tasks in `ML` explain
 218 the scores observed for Pass^k , detailed in Appendix G.

219 This correlation can be used to strategically improve agentic systems. To test this, we augmented the
 220 `SPECTRA` environment with a new tool that works as a hypothesis generator by proposing molecular
 221 fragments for the candidate molecule. While this value might seem insignificant, as Figure 5 shows,

Table 1: **Embedding distance analysis.** Mean cosine distances between consecutive tasks in a chain, between all tasks and tools, and between the minimum or closest tool to each task. The results show that SPECTRA is among the highest values for the measured distances, which could explain the bad results in this environment.

Task	Tools	Mean of consecutive tasks	Verbosity	Distance Tools-Tasks	
				Mean	Mean of Minimum
SPECTRA	15	0.52	brief	0.70	0.52
			workflow	0.75	0.59
			comprehensive	0.84	0.73
MD	4	0.26	brief	0.64	0.58
			workflow	0.65	0.60
			comprehensive	0.72	0.66
OPENCATALYST	6	0.38	brief	0.61	0.38
			workflow	0.60	0.42
			comprehensive	0.71	0.57
ML	12	0.65	brief	0.68	0.39
			workflow	0.65	0.40
			comprehensive	0.72	0.52

we find that with this tool Claude-3.5 Sonnet now performs comparably in both single and chained settings. GPT-4o, which is less proficient in spectroscopic tasks in other benchmarks [13], fails to leverage this additional tool.

4 Conclusion

Large language models have demonstrated remarkable capabilities across diverse tasks, yet significant performance gaps persist in specialized domains like chemistry and materials science. The field has responded with increasing enthusiasm for agentic systems—frameworks that equip LLMs with external tools and reasoning scaffolds. This approach emerged from a compelling belief: that agents could serve as a universal bridge, compensating for any base model’s limitations simply by providing domain-specific tools. The promise was elegant in its simplicity—rather than training specialized models for each domain, we could build flexible agents that adapt to any task through careful framework engineering and tool provisioning.

Our systematic evaluation using our *Corral* framework reveals this foundational assumption to be deeply flawed. Through comprehensive benchmarking across four expert-designed environments spanning molecular dynamics, machine learning, catalysis, and spectroscopy, we demonstrate that the choice of agentic framework—whether ReAct, or tool-calling—plays a surprisingly minor role in determining success. Instead, performance depends critically on the alignment between available tools and task requirements, coupled with the base model’s existing domain knowledge. When models already possess adequate domain expertise, agentic frameworks introduce unnecessary overhead without meaningful benefits. Conversely, when domain knowledge is absent, only meticulously crafted tools that precisely match task demands can bridge the gap—a requirement that eliminates the promised flexibility and demands exactly the kind of deep domain expertise that agents were meant to circumvent.

These findings fundamentally challenge the notion that agentic systems provide a universal solution to model limitations. Rather than offering a flexible path around the need for specialized knowledge, agents reveal themselves to be highly dependent on the very expertise they promised to democratize. Our work highlights an enduring truth: the absence of domain-specific knowledge in foundation models cannot be overcome through framework sophistication alone, but only through targeted tool design that requires the same level of domain expertise as training specialized models—suggesting that domain-specific models are far less obsolete than one might think.

References

- [1] Andrew D. White. "The future of chemistry is language". In: *Nature Reviews Chemistry* 7.7 (May 2023), pp. 457–458. ISSN: 2397-3358. DOI: 10.1038/s41570-023-00502-0. URL: <http://dx.doi.org/10.1038/s41570-023-00502-0>.
- [2] Nawaf Alampara et al. "General purpose models for the chemical sciences". In: *arXiv preprint arXiv: 2507.07456* (2025).
- [3] David Silver and Richard S Sutton. "Welcome to the era of experience". In: *Google AI* 1 (2025).
- [4] John Yang et al. "SWE-agent: Agent-Computer Interfaces Enable Automated Software Engineering". In: *arXiv preprint arXiv: 2405.15793* (2024).
- [5] Bang Liu et al. "Advances and Challenges in Foundation Agents: From Brain-Inspired Intelligence to Evolutionary, Collaborative, and Safe Systems". In: *arXiv preprint arXiv: 2504.01990* (2025).
- [6] Samuel Schmidgall et al. "Agent Laboratory: Using LLM Agents as Research Assistants". In: *arXiv preprint arXiv: 2501.04227* (2025).
- [7] Zhengyao Jiang et al. "AIDE: AI-Driven Exploration in the Space of Code". In: *arXiv preprint arXiv: 2502.13138* (2025).
- [8] Xingyao Wang et al. "OpenHands: An Open Platform for AI Software Developers as Generalist Agents". In: *arXiv preprint arXiv: 2407.16741* (2024).
- [9] Theodore R. Sumers et al. "Cognitive Architectures for Language Agents". In: *arXiv preprint arXiv: 2309.02427* (2023).
- [10] Noah Shinn et al. "Reflexion: Language Agents with Verbal Reinforcement Learning". In: *arXiv preprint arXiv: 2303.11366* (2023).
- [11] Alán Aspuru-Guzik and Varinia Bernales. "The rise of agents: Computational chemistry is ready for (R)evolution". In: *Polyhedron* (July 2025), p. 117707. ISSN: 0277-5387. DOI: 10.1016/j.poly.2025.117707. URL: <http://dx.doi.org/10.1016/j.poly.2025.117707>.
- [12] Santiago Miret and NM Anoop Krishnan. "Enabling large language models for real-world materials discovery". In: *Nature Machine Intelligence* (2025), pp. 1–8.
- [13] Adrian Mirza et al. "A framework for evaluating the chemical knowledge and reasoning abilities of large language models against the expertise of chemists". In: *Nature Chemistry* (2025), pp. 1–8.
- [14] Nawaf Alampara et al. "Probing the limitations of multimodal language models for chemistry and materials research". In: *Nature Computational Science* (2025), pp. 1–10.
- [15] Anthropic. *Model Context Protocol*. Protocol Specification. Open protocol for connecting AI applications to external data sources and tools. Anthropic, PBC, Nov. 2024. URL: <https://modelcontextprotocol.io>.
- [16] Varatheepan Paramanayakam et al. "Less is More: Optimizing Function Calling for LLM Execution on Edge Devices". In: *Design, Automation & Test in Europe Conference, DATE 2025, Lyon, France, March 31 - April 2, 2025*. IEEE, 2025, pp. 1–7. DOI: 10.23919/DATE64628.2025.10992798. URL: <https://doi.org/10.23919/DATE64628.2025.10992798>.
- [17] Guozhao Mo et al. "LiveMCPBench: Can Agents Navigate an Ocean of MCP Tools?" In: *arXiv preprint arXiv: 2508.01780* (2025).
- [18] Mahmoud Mohammadi et al. "Evaluation and Benchmarking of LLM Agents: A Survey". In: *arXiv preprint arXiv: 2507.21504* (2025).
- [19] Nawaf Alampara, Mara Schilling-Wilhelmi, and Kevin Maik Jablonka. "Lessons from the trenches on evaluating machine-learning systems in materials science". In: *arXiv preprint arXiv: 2503.10837* (2025).
- [20] Xuanqi Gao et al. "MCP-RADAR: A Multi-Dimensional Benchmark for Evaluating Tool Use Capabilities in Large Language Models". In: *arXiv preprint arXiv: 2505.16700* (2025).
- [21] Mayk Caldas Ramos, Christopher J. Collison, and Andrew D. White. "A review of large language models and autonomous agents in chemistry". In: *Chemical Science* 16.6 (2025), pp. 2514–2572. ISSN: 2041-6539. DOI: 10.1039/d4sc03921a. URL: <http://dx.doi.org/10.1039/D4SC03921A>.

- 306 [22] Siddharth Narayanan et al. “Aviary: training language agents on challenging scientific tasks”.
307 In: *arXiv preprint arXiv: 2412.21154* (2024).
- 308 [23] Thomas Carta et al. “Grounding Large Language Models in Interactive Environments with
309 Online Reinforcement Learning”. In: *arXiv preprint arXiv: 2302.02662* (2023).
- 310 [24] Muning Wen et al. “Reinforcing language agents via policy optimization with action decompo-
311 sition”. In: *arXiv preprint arXiv:2405.15821* (2024).
- 312 [25] Yifan Song et al. “Trial and error: Exploration-based trajectory optimization for llm agents”.
313 In: *arXiv preprint arXiv:2403.02502* (2024).
- 314 [26] Shunyu Yao et al. “ReAct: Synergizing Reasoning and Acting in Language Models”. In:
315 *International Conference on Learning Representations* (2022).
- 316 [27] Chan Hee Song et al. “LLM-Planner: Few-Shot Grounded Planning for Embodied Agents
317 with Large Language Models”. In: *arXiv preprint arXiv: 2212.04088* (2022).
- 318 [28] Matthew K Horton et al. “Accelerated data-driven materials science with the Materials Project”.
319 In: *Nature Materials* (2025), pp. 1–11.
- 320 [29] Yushi Bai et al. “LongBench v2: Towards Deeper Understanding and Reasoning on Realistic
321 Long-context Multitasks”. In: *arXiv preprint arXiv: 2412.15204* (2024).
- 322 [30] Cheng-Ping Hsieh et al. “RULER: What’s the Real Context Size of Your Long-Context
323 Language Models?” In: *arXiv preprint arXiv: 2404.06654* (2024).

Appendix

A Related Work

LLM-based Agents LLM-based agents demonstrated proficiency in coding tasks [1]. This has extended to specialized domains including ML development [2], algorithm optimization [3], and AI research automation [4]. *CodeAct Agent* is a multi-agent system that enables agents to either converse with humans in natural language for clarification and confirmation or perform tasks by executing code, such as running bash commands, Python scripts, or browser-specific programming languages [5]. Similarly, *AIDE* is an ML-focused agent that autonomously drafts, debugs, and optimizes code through iterative tree-search strategies based on user-defined goals and metrics [6].

In science, agents have targeted distinct phases of the research pipeline, including data retrieval [7, 8], simulation execution [9, 10, 11, 12], and data analysis [13]. Some notable examples are those that can run certain lab experiments [14, 15, 16].

Scientific Agent Benchmarks The advancements in agents are facilitated by high-quality benchmarking frameworks that rigorously quantify agent capabilities in coding [17, 18, 19, 20], chemistry [21], and materials science [22, 23]. Tian et al. [24] demonstrated limitations of agents replicating real-world scientific scripts. Even top-performing models achieved below 10 % accuracy in generating specialized scientific code. This performance gap extends to other scientific contexts, as shown by Cui et al. [25]’s multi-domain benchmark featuring complex tasks like DFT calculations and materials property extraction. Complementary approaches like *Aviary* [26] or *AutoGen* [27] address system design by isolating environments from agents. These “agent gym” frameworks enable tool decoupling and targeted optimization, yielding performance gains.

Existing benchmarks target specific scientific tasks given diverse evaluation needs [28, 29]. There have also been attempts to measure whether the provided tools are truly relevant or if the agent understands when and how to use them. One proposed way to assess this is by checking if expected tools are used at each step [30], but this approach is inadequate for environments requiring open-ended planning and strategic reasoning [28, 19]. In such cases, the agent’s ability to decompose a task and strategically select tools becomes paramount, and currently, there is no way to measure this ability.

Overall, there is no systematic understanding that can be used as practical guidance for improving scientific agents. By systematically ablating the components within the agentic framework for a set of scientific environments, our work identifies critical performance limitations and provides actionable pathways for agent improvement.

B Metrics

To evaluate the overall performance of the benchmark, we employ a suite of metrics that capture different aspects of the agent’s capabilities, from correctness and efficiency to resource consumption.

B.1 Average Score

The average score provides a measure of the mean performance across all tasks. It is calculated as the average of the mean scores for each individual task.

Let T be the set of all task IDs, and for each task $t \in T$, let S_t be the set of scores for each trial of that task. The average score for a single task t is given by:

$$\mu_t = \frac{1}{|S_t|} \sum_{s \in S_t} s \quad (\text{B1})$$

The overall average score is then the mean of these individual task averages:

$$\text{Average Score} = \frac{1}{|T|} \sum_{t \in T} \mu_t \quad (\text{B2})$$

365 B.2 Overall Success Rate

366 The overall success rate is the average of the success rates of all individual tasks. A trial is considered
367 successful if its score is greater than zero and no error message is generated.

368 Let T be the set of all tasks. The success rate for a single task t , denoted as SR_t , is the proportion of
369 successful trials for that task. The overall success rate is then:

$$\text{Overall Success Rate} = \frac{1}{|T|} \sum_{t \in T} SR_t \quad (\text{B3})$$

370 B.3 Pass@ k

371 The pass@ k [31] metric evaluates the probability that at least one out of k trials for a given task
372 succeeds. It is a valuable metric for assessing the likelihood of achieving a correct solution within a
373 limited number of attempts. The overall pass@ k is the average of the task-specific pass@ k values.

374 For a single task, let n be the total number of trials and c be the number of successful trials. The
375 probability of a single trial succeeding is $p = c/n$. The pass@ k is then calculated as:

$$\text{pass@}k = 1 - (1 - p)^k = 1 - \left(1 - \frac{c}{n}\right)^k \quad (\text{B4})$$

376 The overall pass@ k is the average of this value across all tasks:

$$\text{Overall pass@}k = \frac{1}{|T|} \sum_{t \in T} \text{pass@}k_t \quad (\text{B5})$$

377 B.4 Pass $^{\wedge}k$

378 The pass $^{\wedge}k$ (pass-hat- k) metric [31] measures the probability that all k trials for a given task succeed.
379 This metric is a stricter evaluation of reliability and consistency. The overall pass $^{\wedge}k$ is the average of
380 the task-specific pass $^{\wedge}k$ values.

381 Using the same notation as for pass@ k , the pass $^{\wedge}k$ for a single task is given by:

$$\text{pass}^{\wedge}k = p^k = \left(\frac{c}{n}\right)^k \quad (\text{B6})$$

382 The overall pass $^{\wedge}k$ is the average of this value across all tasks:

$$\text{Overall pass}^{\wedge}k = \frac{1}{|T|} \sum_{t \in T} \text{pass}^{\wedge}k_t \quad (\text{B7})$$

383 B.5 Token Usage

384 To measure the computational cost, we record the token usage for each trial. The total token usage is
385 the sum of tokens consumed across all trials in the benchmark. This can be broken down into prompt
386 tokens, completion tokens, and total tokens.

387 B.6 Tool Call Statistics

388 We also track the usage and success of tool calls made by the agent. This includes the total number
389 of successful and failed tool calls across all trials and tasks. For failed tool calls, the failure is
390 also classified to `invalid_tool` (invalid tool name), `invalid_args` (invalid tool arguments) or
391 `execution_error` (error in executing the tool).

392 B.7 Duration Metrics

393 We measure various aspects of the benchmark’s execution time:

- 394 • **Total Tool Execution Duration:** The sum of the execution time for all tool calls across all
395 trials.
- 396 • **Overall Average Trial Duration:** The average duration of a single trial across all tasks.
- 397 • **Total Duration:** The total time taken to run all trials for all tasks in the benchmark.




Figure C.1: **Illustration of the *Corral* framework.** The architecture is composed of the *Corral* server (implementing the environment dynamics), agents (the policies π), and the *Corral* runner (orchestrating the agent-environment interaction over a series of tasks).

To facilitate interaction between the agent and the environments, we have developed a benchmark server with a *RESTful* API. This server exposes a set of endpoints that allow the agent to retrieve task information, interact with the provided tools, and submit its final answer. The API is designed to be clear, consistent, and stateless, adhering to the principles of REST. The available endpoints are detailed in Table C.1. Each endpoint is described with its HTTP method, URL path, a brief description of its functionality, and any relevant path or query parameters.

Table C.1: **Benchmark Server API Endpoints.** The table specifies all the endpoints that the server hosting the environments admits.

Method	Path	Description
Task and Environment Discovery		
GET	/tasks	Retrieves a list of all available task IDs in the benchmark.
GET	/dependency_chain	Indicates whether any tasks in the benchmark are part of a dependency chain.
GET	/tasks/{task_id}/prompt	Fetches the specific prompt for a given task. The {task_id} is a path parameter representing the unique identifier for the task.
GET	/tasks/{task_id}/guide	Returns the complete environment guide for a task, including the prompt and a detailed description of the available tools. Accepts an optional verbosity query parameter (e.g., MINIMAL or WORKFLOW) to control the level of detail in the tool descriptions.
Tool Interaction and Introspection		
GET	/tasks/{task_id}/tools	Returns a structured JSON object detailing the available tools for a task, compatible with function-calling APIs. The level of detail in the argument descriptions is controlled by the optional verbosity query parameter.

Table C.1 continued from previous page

Method	Path	Description
GET	/tasks/{task_id}/tools/guide	Provides a guide to the tools available for a specific task. Also accepts the optional verbosity query parameter.
POST	/tasks/{task_id}/tools/execute	Executes a specified tool with the provided arguments within the task’s environment. The tool name and arguments are sent in the request body.
State Management and Task Submission		
GET	/tasks/{task_id}/state	Retrieves the current state of the specified task environment.
POST	/tasks/{task_id}/submit	Submits a final answer for the task. The submitted answer is in the request body. The response includes the score and the final state of the trial.
GET	/tasks/{task_id}/status	Gets the completion status of a task, including whether it is completed, the final score, the submitted answer, and tool usage statistics.
GET	/tasks/{task_id}/trials	Retrieves the states of all trials that have been run for a specific task.
GET	/tasks/{task_id}/trials/{trial_id}	Fetches the state of a specific trial for a given task, identified by its unique {trial_id}.

D Structured Tool Descriptions and Verbosity Levels

To systematically ablate the information provided to the agent, we developed a structured, tag-based format for all tool docstrings. This allows us to programmatically control the content of the tool descriptions presented to the LLM at different verbosity levels. When a new tool is developed, its function docstring is written once in a comprehensive, structured format. During the setup of a benchmark run, our framework dynamically parses this complete docstring. Based on the selected verbosity strategy (e.g., BRIEF, WORKFLOW), it then filters and formats the content to generate the precise, tailored tool description that will be exposed to the agent.

D.1 Documentation Tags

Each part of the tool’s documentation is enclosed in a specific tag, allowing for granular filtering. The primary tags used are:

[BRIEF] A single-sentence, high-level summary of the tool’s function.

[DETAILED] A more thorough explanation of what the tool does and how it works internally.

[PROCEDURAL] Guidance on when and why to use the tool in a problem-solving process.

[CONTEXTUAL] Information about the tool’s operational context, such as its dependencies or data sources.

[WORKFLOW_INTEGRATION] An explicit, structured guide on how the tool fits into a multi-step plan, outlining prerequisites, the current action, and follow-up steps.

[SYNTACTICAL] Concrete, valid usage examples of the tool’s syntax.

[RAISES] A list of potential exceptions or errors the tool might raise, including conditions and recovery suggestions.

[LIMITATIONS] Known limitations or edge cases where the tool might not perform as expected.

[EXAMPLES] Illustrative examples, often for arguments or return values.

428 D.2 Verbosity Levels

429 Our ablation studies focus on three distinct verbosity levels, which are constructed by including
430 content from a specific subset of the tags described above.

- 431 • **BRIEF:** Includes only content from the [BRIEF] tags. This provides the agent with the most
432 minimal functional description.
- 433 • **WORKFLOW:** Includes content from [BRIEF], [DETAILED], [PROCEDURAL],
434 [CONTEXTUAL], and [WORKFLOW_INTEGRATION]. This level is designed to give the
435 agent rich strategic and operational context.
- 436 • **COMPREHENSIVE:** Includes all available tags. This level equips the agent with complete
437 knowledge for robust execution, including syntax, error handling, and limitations.

438 D.3 Example: The `get_structure_from_mp_text` Tool

439 The following examples demonstrate the exact text shown to the agent for the same tool at each of the
440 three verbosity levels. Note that the descriptive tags themselves are filtered out, and only the enclosed
441 content is presented.

Example: BRIEF Verbosity Description

Retrieve a pymatgen structure from Materials Project using its API and return CIF content as text.

Args:

`mp_id`: Materials Project identifier string.

Returns:

`str`: CIF content string containing the crystal structure data.

442

Example: WORKFLOW Verbosity Description

Retrieve a pymatgen structure from Materials Project using its API and return CIF content as text.

This tool connects to the Materials Project database to download crystal structure data for a given material ID. It retrieves the structure object and converts it to CIF (Crystallographic Information File) format, which is the standard format for storing crystal structure information. CIF is then returned as string

When to use: Use when you need to retrieve bulk crystal structures from the Materials Project database. Best suited for materials with known MP IDs. Usually first step in simulation workflows. Recommended for obtaining crystal structures for preparing bulk structures, supercells, bulk cells, slabs etc. Avoid when you need multiple structures.

How this tool works: Connects to Materials Project API using authentication key. Searches for the specified material ID (MP ID) in the database (MP ID is given as input parameter or if other tools are available to search for MP ID based on available information, then use those tools). Retrieves the pymatgen Structure object containing atomic positions and lattice parameters. Converts the structure to CIF format string for compatibility with other tools. Returns standardized crystallographic data suitable for further processing.

Workflow: 1. [PREREQUISITE] Ensure that the other more specific tools are not suitable and you don't have to retrieve multiple structures 2. [CURRENT] Apply this tool with a valid MP ID to retrieve bulk structure 3. [FOLLOW_UP] Use the CIF output with slab generation tools like `enumerate_slabs_text` to create slab structures

443

Args:

mp_id: Materials Project identifier string. The unique identifier used by Materials Project to catalog materials. Should be in the format "mp-XXXXX" where XXXXX is a numerical ID. This ID corresponds to a specific material entry in the Materials Project database.

Returns:

str: CIF content string containing the crystal structure data. A properly formatted CIF (Crystallographic Information File) string containing all necessary information about the crystal structure, including lattice parameters, atomic positions, space group, and symmetry operations. This format is widely compatible with crystallographic software and other structure analysis tools.

444

Example: COMPREHENSIVE Verbosity Description

Retrieve a pymatgen structure from Materials Project using its API and return CIF content as text.

This tool connects to the Materials Project database to download crystal structure data for a given material ID. It retrieves the structure object and converts it to CIF (Crystallographic Information File) format, which is the standard format for storing crystal structure information. CIF is then returned as string.

When to use: Use when you need to retrieve bulk crystal structures from the Materials Project database. Best suited for materials with known MP IDs. Usually first step in simulation workflows. Recommended for obtaining crystal structures for preparing bulk structures, supercells, bulk cells, slabs etc. Avoid when you need multiple structures.

How this tool works: Connects to Materials Project API using authentication key. Searches for the specified material ID (MP ID) in the database...

Workflow: 1. [PREREQUISITE] Ensure that the other more specific tools are not suitable... 2. [CURRENT] Apply this tool with a valid MP ID to retrieve bulk structure 3. [FOLLOW_UP] Use the CIF output with slab generation tools...

Usage examples:

```
get_structure_from_mp_text("mp-149") # Silicon structure
get_structure_from_mp_text("mp-20066") # CO2 structure
```

Exceptions:

- **ConnectionError:** When unable to connect to Materials Project API. Recovery: Check internet connection and MP_API_KEY.
- **KeyError:** When the specified MP ID is not found. Recovery: Verify MP ID syntax and existence on the Materials Project website.

Limitations:

- Limited to materials available in the Materials Project database.
- May not include the most recent experimental structures.

Args:

mp_id: Materials Project identifier string. The unique identifier used by Materials Project... Should be in the format "mp-XXXXX"... Syntax: "mp-" followed by digits (e.g., "mp-149", "mp-20066"). Examples: "mp-149" (Silicon), "mp-20066" (CO2).

445

Returns:

str: CIF content string containing the crystal structure data. A properly formatted CIF string containing all necessary information... Example:

```
# generated using pymatgen
data_Si
_symmetry_space_group_name_H-M    'P 1'
_cell_length_a    3.83996459
...
```

446

447 E Corral Framework Usage Examples

448 The *Corral* framework is built using such a modular design that is simple to implement agents,
449 add new environments, and run the benchmark. For example, the code block below shows how to
450 benchmark an agent using the *Corral* framework. It needs to set up the environment, create a logger
451 for tracking experiments, and then run a benchmark with an agent, outputting the results such as the
452 overall score.

```
453 from corral.evaluate import BenchmarkInterface, MatAgentBenchmark
454 from corral.agents.react import ReActAgent
455 from corral.report import CorralWandbLogger
456
457 # Setup environments server running
458 interface = BenchmarkInterface("http://localhost:8000")
459
460 # Setup the WandB logger
461 wandblogger = CorralWandbLogger(
462     project="corral",
463     group="experiment_group",
464     name="run_name",
465 )
466
467 # Setup the agent
468 agent = ReActAgent(model="gpt-4o", max_iterations=10, temperature=0.1)
469
470 # Run benchmark
471 runner = MatAgentBenchmark(interface, agent, logger=wandblogger)
472 result = runner.bench()
473
474 print(f"Overall score: {result.total_score:.2f}")
475
476
```

478 The next code block focuses on setting up and running a *Corral* environment server. In it, a custom
479 environment is defined by implementing the *Environment* class, specifying the task, problem, and
480 expected answer. The server is then created to host the benchmark tasks and make them accessible to
481 agents for evaluation.

```
482 from corral.base import Environment
483 from corral.server import create_benchmark_server
484
485 class MyEnvironment(Environment):
486     def __init__(self, task_id: str, problem: str, answer: str):
487         self.problem = problem
488         self.correct_answer = answer
489         super().__init__(task_id)
```

```

492
493     # Add your tools
494     self.add_tool(my_custom_tool)
495
496     def get_task_prompt(self) -> str:
497         return f"Solve this problem: {self.problem}"
498
499     def score(self) -> float:
500         if self.state.submitted_answer is None:
501             return 0.0
502         return 1.0 if self.state.submitted_answer == self.
503             correct_answer else 0.0
504
505
506     # Define your tasks
507     environments = {
508         "task_1": MyEnvironment("task_1", "Problem 1", "Answer 1"),
509         "task_2": MyEnvironment("task_2", "Problem 2", "Answer 2"),
510     }
511
512     # Create server
513     if __name__ == "__main__":
514         app = create_benchmark_server(environments)
515         import uvicorn
516
517         uvicorn.run(app, host="0.0.0.0", port=8000)

```

519 Similarly, the code block below demonstrates how to create a custom agent within the *Corral*
520 framework. To that, one can define a custom agent class that extends `BaseAgent` and implement the
521 logic for the agent's task processing. This agent interacts with the benchmark interface to retrieve
522 task information and use tools to solve problems.

```

523
524 from corral.agents.base_agent import BaseAgent
525 from corral.evaluate import BenchmarkInterface
526
527
528 class MyAgent(BaseAgent):
529     def __init__(self, model: str, **kwargs):
530         super().__init__(model, **kwargs)
531         # Add your agent-specific initialization
532
533     def run(self, interface: BenchmarkInterface, task_id: str) -> str:
534         # Get task information
535         guide = interface.get_task_guide(task_id)
536
537         # Your agent logic here
538         # Use interface.execute_tool() to call tools
539
540         return "Your final answer"

```

542 Lastly, the next code block defines a custom tool for the environment,
543 `calculate_molecular_weight`, which can be used by agents to perform specific tasks, in
544 this case, calculating the molecular weight of a given chemical formula.

```

545
546 from corral.utils import tool
547
548
549 @tool
550 def calculate_molecular_weight(formula: str) -> float:
551     """[BRIEF] Calculate molecular weight from chemical formula. [/
552     BRIEF]

```

```

553
554     Args:
555         formula: Chemical formula (e.g., 'H2O', 'CH4')
556
557     Returns:
558         Molecular weight in g/mol
559     """
560     # Implementation here
561     pass
562

```

563 F Detailed Environment Descriptions

564 We compiled four expert-curated environments, with over 35 domain-specific tools, and over 20
565 specialized scoring functions.

566 F.1 Overview of the Environments

567 We implement multiple environments, covering diverse problems in chemistry and materials research.

568 **Spectra** Spectroscopic analysis is a fundamental technique in chemistry for determining the
569 structure of unknown compounds, using data from methods such as nuclear magnetic resonance
570 (NMR), infrared (IR) spectroscopy, or mass spectrometry (MS). This environment evaluates LLM
571 agents on the task of molecular structure elucidation from spectroscopic data, requiring them to plan
572 the experiments and output a valid SMILES representation of the compounds. This environment
573 mimics a virtual lab where the agent has to first plan what experiments to run, then propose different
574 molecular hypotheses based on the different experiments that it performs, validate those hypotheses,
575 and finally return the valid elucidated molecular structure.

576 **MD** MD simulations are widely used to study the atomic-scale behavior of materials, providing
577 information on structural, thermal, and mechanical properties at timescales that are difficult to access
578 experimentally [32]. The MD environment explores the ability of agents to use this technique with
579 the software LAMMPS [33]. Agents in this environment can retrieve atomic structures from MP [34],
580 write LAMMPS input scripts, and perform simulations to compute target properties. While the
581 implementation is shared, we define three separate environments based on the task being performed:
582 *Surface Energy*, *Melting*, and *Quenching* simulations.

583 **OpenCatalyst** This environment is inspired by the *Open Catalyst Project* computational work-
584 flow [35]. Catalysis plays a pivotal role in addressing critical global challenges, including renewable
585 energy generation (e.g., fuel cells, solar fuels) and sustainable chemical production (e.g., ammonia
586 synthesis). The discovery of novel catalysts is challenging due to the vast chemical space and diversity
587 of possible surfaces, terminations, adsorption sites, and adsorbate orientations. This environment
588 evaluates agents on performing the key and early steps of a typical catalysis simulation pipeline,
589 which is to create a catalyst slab-adsorbate configuration.

590 **ML** ML models are increasingly used in materials science to accelerate property predictions and
591 reduce reliance on costly first-principles simulations or experiments. This environment evaluates
592 agents on constructing and training an XGBoost model [36] to predict the formation energies of
593 oxide, nitride, and sulphide polymorphs using data from MP. Agents must retrieve relevant crystal
594 structures, curate polymorph datasets, perform feature engineering, train, and evaluate the model with
595 both test set metrics and cross-validation, and save the results in a structured format.

596 Below, we provide detailed descriptions of each environment, including associated tasks and scoring
597 functions used to evaluate agent performance.

598 F.2 Spectra

599 This environment simulates a virtual lab, providing tools to run various experimental analyses, such
600 as carbon or proton NMR. The scoring function of the single task is simply comparing the smiles to

601 check if they correspond to the same molecule. The different scoring functions of the chained tasks
602 correspond to counting the number of certain atoms, some functional groups, etc.

603 F.2.1 Tools

604 Note that for the experimental tools in the SPECTRA environment, no arguments are required. This is
605 because agents are prompted in an environment in which the sample at hand is already loaded into
606 the experimental equipment. For the actual "experimental" tools, we used simulation tools available
607 through API services [37, 38, 39, 40]

Table F.2: Spectra **Tools**. The table describes each tool available in the spectra environment, including its arguments, return values, and potential errors.

Tool
<code>get_formula_from_smiles</code>
Arguments: <code>smiles</code> (string): SMILES representation of a molecule.
Return: Generates the chemical formula of the molecule in Hill notation (C, H, then alphabetical order). Uses RDKit to parse the SMILES string and compute the formula. Returns a string formula if successful, or an error message if the SMILES is invalid.
Raise: Raises Exception for unparseable SMILES strings.
<code>search_by_smiles</code>
Arguments: <code>smiles</code> (string): SMILES representation of the compound to search for in the NMRShift database.
<code>top_k</code> (int, optional, default 10): Maximum number of results to return.
Return: Searches the NMRShift database for entries matching or chemically similar to the provided SMILES. Returns a list of dictionaries containing relevant information about each compound, such as SMILES, chemical shifts, and other properties. Results are sorted by similarity score in descending order.
Raise: Raises Exception if the search fails or if <code>top_k</code> is invalid.
<code>retrieve_aromatic_protons_shifts</code>
Arguments: None
Return: Retrieves the proton chemical shift ranges for aromatic hydrocarbons. Returns a string representation of a list of dictionaries, each containing the type of proton (e.g., aldehyde, aromatic, alkene) and its corresponding chemical shift range in ppm relative to TMS.
Raise: No exceptions are raised.
<code>retrieve_carbon_shifts</code>
Arguments: None
Return: Retrieves the carbon chemical shift ranges for various functional groups in organic compounds. Returns a string representation of a list of dictionaries, each containing the functional group (e.g., CH_3- , $\text{R}_3\text{C}-$, $=\text{CH}$, Ketones, Aldehydes) and its corresponding chemical shift range in ppm relative to TMS.
Raise: No exceptions are raised.

Tool
carbon_nmr_spectra Arguments: None Return: Measures the ^{13}C NMR spectra for the given sample using the <code>get_c13_nmr_prediction</code> function. Returns the ^{13}C NMR spectra as a string in ACS-inspired publication format. Raise: Raises an exception if an error occurs during measurement.
proton_nmr_spectra Arguments: None Return: Returns the ^1H NMR spectra for the given compound using the <code>get_h_nmr_prediction</code> function. Provides the ^1H NMR spectra as a string following ACS-inspired publication format. Raise: Raises an exception if an error occurs during measurement.
ir_spectra Arguments: None Return: Returns the IR spectra for the given compound using the <code>get_ir_prediction</code> function. Provides the IR spectra as a string following standard IR notation. Raise: Raises an exception if an error occurs during measurement.
hsqc_nmr_spectra Arguments: None Return: Returns the HSQC (Heteronuclear Single Quantum Coherence) NMR spectra for the given compound by making an external API call and formatting the result in standard NMR notation. Raise: Raises an exception if an error occurs during measurement or if the HSQC spectrum is unavailable.
mass_spectrometry_spectra Arguments: None Return: Returns the mass spectrometry spectra for the given compound using Electrospray Ionization (ESI) by making an external API call. Outputs the spectrum in the format <code>m/z value (intensity), ...</code>. Raise: Raises an exception if an error occurs during measurement or if the SMILES string is invalid.
retrieve_isotope_distribution Arguments: None Return: Returns a predefined dictionary containing the isotopic distribution of common elements in organic chemistry (C, H, S, Cl, Br, I, F, N, O). Raise: No calculations or API calls are performed; values are typical for organic compounds.

Tool
<code>retrieve_dbe_formula</code> Arguments: None Return: Returns the formula for calculating the Double Bond Equivalent (DBE), also known as Degree of Unsaturation (DU) or Index of Hydrogen Deficiency (IHD). Raise: No exceptions are raised.
<code>obtain_isomers</code> Arguments: <code>smiles</code> (string): SMILES representation of the compound. Return: Retrieves structural isomers for a given compound using its SMILES representation. Uses a remote function to query the PubChem database and return a list of SMILES strings for compounds with the same molecular formula. Raise: Returns an empty list if no isomers are found.
<code>obtain_isomers_from_molecular_formula</code> Arguments: <code>molecular_formula</code> (string): Molecular formula of the compound. Return: Retrieves structural isomers for a given molecular formula. Uses a remote function to query the PubChem database and return a list of SMILES strings for compounds with the same molecular formula. Raise: Returns an empty list if no isomers are found.
<code>validate_smiles</code> Arguments: <code>smiles</code> (string): SMILES representation of the compound. Return: Validates a SMILES string to check if it represents a valid chemical structure. Returns True if the conversion is successful (valid SMILES) and False otherwise. Raise: No exceptions are raised.

608 F.2.2 Scoring Functions

609 The scoring functions for the SPECTRA environment are described in Table F.3. Note that the
610 `score_molecule` function is the one used to score the single task.

Table F.3: Spectra **scoring functions**. The table describes each scoring function, including its arguments, return values, and behavior.

Function
<code>score_molecule</code> Arguments: <code>prediction</code> (str, SMILES): Predicted SMILES string. <code>ground_truth</code> (str, SMILES): Ground truth SMILES string. Return: Returns 1.0 if the canonical SMILES match, ignoring stereochemistry, otherwise 0.0.
<code>score_molecule_fragments</code>

Function
Arguments: prediction (list[str] or str): Predicted fragments (list or string). ground_truth (str, SMILES): Ground truth SMILES string. Return: Returns 1.0 if all predicted fragments are substructures of the ground truth molecule, otherwise 0.0.
score_formula_match Arguments: prediction (str, formula): Predicted molecular formula. ground_truth (str, SMILES): Ground truth SMILES string. Return: Returns 1.0 if predicted molecular formula matches the ground truth formula, otherwise 0.0.
validate_dbe_consistency Arguments: prediction (int): Predicted degree of unsaturation (DBE). ground_truth (str, SMILES): Ground truth SMILES string. Return: Returns 1.0 if predicted DBE matches calculated DBE from SMILES, otherwise 0.0.
score_isotopic_distribution Arguments: prediction (list[str], element symbols): Predicted isotopic distribution of elements. ground_truth (str, SMILES): Ground truth SMILES string. Return: Returns 1.0 if predicted elements with isotopic distributions match those in the molecule, otherwise 0.0.
score_num_hydrogen_symmetry_classes Arguments: prediction (int): Predicted number of hydrogen symmetry classes. ground_truth (str, SMILES): Ground truth SMILES string. Return: Counts unique hydrogen symmetry classes using canonical ranking. Returns 1.0 if predicted count matches, otherwise 0.0.
score_num_carbon_symmetry_classes Arguments: prediction (int): Predicted number of carbon symmetry classes. ground_truth (str, SMILES): Ground truth SMILES string. Return: Counts unique carbon symmetry classes using canonical ranking. Returns 1.0 if predicted count matches, otherwise 0.0.
score_num_aromatic_carbons Arguments: prediction (int): Predicted number of aromatic carbons. ground_truth (str, SMILES): Ground truth SMILES string. Return: Counts aromatic carbons in the molecule. Returns 1.0 if predicted count matches, otherwise 0.0.

Function

score_num_ch3_groups

Arguments: prediction (int): Predicted number of CH₃ (methyl) groups.
ground_truth (str, SMILES): Ground truth SMILES string.

Return: Counts CH₃ groups (methyl) in the molecule. Returns 1.0 if predicted count matches, otherwise 0.0.

score_num_carbonyl_groups

Arguments: prediction (int): Predicted number of carbonyl groups.
ground_truth (str, SMILES): Ground truth SMILES string.

Return: Counts carbonyl groups (C=O) in the canonical tautomer of the molecule. Returns 1.0 if predicted count matches, otherwise 0.0.

611 **F.2.3 Tasks and Subtasks**

612 Here, we present the prompts for a single task along with its associated subtasks. The structure and
613 format are consistent across other tasks within the SPECTRA environment. Note that for these tasks,
614 there are no input parameters since the different machines are already configured.

Organic Compound Analysis

Analyze the provided organic compound sample in a lab environment and output the SMILES string, while minimizing resource consumption due to the costly nature of the process.

Subtask 1: Molecular Formula

State the molecular formula for the compound analyzed.

Subtask 2: Double Bond Equivalents (DBE)

Calculate the Double Bond Equivalents (DBE) for the compound analyzed. Provide your answer as an integer.

Subtask 3: Identify Elements from Mass Spectrum

Based on the isotope distribution patterns visible in the mass spectrum, list all elements that can be definitively identified. Format your answer as a list of element symbols: [element1; element2; element3].

Subtask 4: Chemically Equivalent Carbon Atoms

How many different types of chemically equivalent carbon atoms are present in this compound? (Count the number of unique carbon environments, not the total number of carbons)

615

Subtask 5: Chemically Equivalent Hydrogen Atoms How many different types of chemically equivalent hydrogen atoms are present in this compound? (Count the number of unique hydrogen environments, not the total number of hydrogens)
Subtask 6: Aromatic Carbons Count State the count of aromatic carbons in the molecule.
Subtask 7: Methyl Groups Count How many methyl (CH₃) groups are present in the sample's molecule?
Subtask 8: Carbonyl Groups Count Indicate the number of C=O groups present in the molecule.
Subtask 9: SMILES for Connected Fragments Return the SMILES strings for the connected fragments of the compound.
Subtask 10: Complete SMILES Provide the SMILES string representing the compound in the sample.

616

617 F.3 MD

618 We deploy the LAMMPS package [33] using the Modal Inference serverless infrastructure [41]. All
 619 input/output operations and tool executions are integrated with the Modal platform. This gives the
 620 flexibility to execute LAMMPS simulations without requiring local or cloud-based installation, and it
 621 automatically provides computational resources at the time of running the simulation. Structure files
 622 and potential files required for MD simulations are stored in separate volumes, which the agent can
 623 access directly during execution. Table F.4 lists the types of potentials available to the agent and the
 624 corresponding materials.

Table F.4: Potentials available to the agent.

Potential Type	Applicable Materials
Embedded Atom Method (EAM) ReaxFF	Aluminum, Magnesium, Copper Silicon, Alpha-Quartz, Wollastonite

625 This environment provides the underlying implementation for three task environments: Surface
 626 Energy, Melting Dynamics, and Quenching Dynamics.

627 F.3.1 Tools

628 Table F.5 provides a detailed description of the tools used in the MD environment, including what their
 629 arguments, return objects, and the possible exceptions raised.

Table F.5: **Tools for the MD environment.** The table describes each tool available in this environment, including its arguments, return values, and potential errors.

Tool
<code>get_potential_metadata</code> Arguments: <code>file_path</code> (string): Path to the LAMMPS potential file. Return: Returns structured metadata: potential type, supported elements, pair style. Raise: Raises <code>ValueError</code> if file is unrecognized or empty.
<code>get_structure_from_mp_text</code> Arguments: <code>mp_id</code> (string): Materials Project ID. <code>file_path</code> (string): Destination path for the CIF file. Return/Behavior: Retrieves crystal structure from Materials Project, saves CIF file, and returns a confirmation message. Raise: Raises <code>Exception</code> if retrieval or saving fails.
<code>convert_structure_to_lammps_data</code> Arguments: <code>structure_path</code> (string): Path to the CIF file. <code>output_file</code> (string): Output LAMMPS data path. <code>atom_style</code> (string, optional, default “charge”): Atom style for LAMMPS conversion. Return/Behavior: Converts CIF to a LAMMPS <code>.data</code> file using the specified atom style. Returns a confirmation message. Raise: Raises <code>Exception</code> if the files are invalid or the atom style is unsupported.
<code>run_lammps</code> Arguments: <code>input_file</code> (string): Path to the LAMMPS input script. Return/Behavior: Runs LAMMPS simulation via a remote backend, generates a <code>.log</code> file, and returns a success message with the log file path. Raise: Raises <code>ValueError</code> for simulation errors or a generic <code>Exception</code> for unexpected issues.

630 F.3.2 Scoring Functions

631 For each task environment, the ground truth is generated using predefined simulation parameters.
632 Agents are then evaluated by comparing their outputs against this ground truth, given the same
633 simulation parameters. For numerical outputs, a tolerance of 2% is applied to account for stochastic
634 variations in molecular dynamics (MD) simulations.

Table F.6: **Scoring Functions in the MD environment.** The table describes each scoring function available in this environment, including its arguments, return values, and behavior.

Function
<code>check_potential_file</code>

Function

Arguments: `target` (string): Path to the reference potential file.

Return: Returns a scoring function. Checks if the result filename matches `target` and verifies file existence via Modal. Returns 1.0 if both pass, otherwise 0.0.

`check_numerical`

Arguments: `target` (float): Expected numerical value.

`tolerance` (float): Allowed relative tolerance.

Return: Returns a scoring function. Parses the numeric result (string, JSON, or dict), compares against `target ± tolerance`, and checks associated file existence via Modal. Returns 1.0 if all checks pass, otherwise 0.0.

`check_structure`

Arguments: `target` (string): Path to the reference LAMMPS .data file.

`atom_style` (string): LAMMPS atom style.

Return: Returns a scoring function. Reads remote LAMMPS .data result via Modal, loads reference and result structures with pymatgen's `LammpsData`, compares using `StructureMatcher`. Returns 1.0 if structures match, otherwise 0.0. Temporary files are automatically cleaned.

635 F.3.3 Tasks and Subtasks

636 Here we provide an example of a task and its associated subtask prompts for Aluminum in each task
637 environment. The prompt structure is similar for other materials.

Surface Energy of Aluminum

Calculate the surface energy of Aluminum for the given surface by removing periodic boundary conditions in the given direction and computing the energy difference per unit area between the relaxed bulk and the relaxed slab configuration. Use the Embedded Atom Method (EAM) potential and the following simulation parameters. The simulation should replicate the unit cell 5 times in each of the x, y, and z directions. Return the surface energy in electronvolts/Å² units.

Input Parameters:

```
{
  "units": "metal",
  "minimiser": "conjugate gradient",
  "energy_tolerance": 1e-10,
  "force_tolerance": 1e-10,
  "maximum_iterations": 100000.0,
  "maximum_evaluations": 100000.0,
  "materials_project_id": "mp-134",
  "direction": [0, 0, 1],
  "atom_style": "atomic"
}
```

638

Subtask 1 : Retrieve Aluminum Structure

Description: Retrieve structure of Aluminum and convert it to LAMMPS data format. Return the absolute path of the final structure.

Input Parameters:

```
{
  "materials_project_id": "mp-134",
  "atom_style": "atomic"
}
```

Subtask 2 : Procure EAM Potential File

Description: Procure the appropriate Embedded Atom Method (EAM) potential file for Aluminum for MD simulations. Return the absolute path to the potential file.

Input Parameters: None

Subtask 3 : Bulk Relaxation for Surface Energy

Description: Given paths to LAMMPS-compatible structure file and EAM potential file, perform structure relaxation of Aluminum (5×5×5 replication). Report total energy (eV) and saved path.

Input Parameters:

```
{
  "units": "metal",
  "pbc": "true",
  "minimiser": "conjugate gradient",
  "energy_tolerance": 1e-10,
  "force_tolerance": 1e-10,
  "maximum_iterations": 1e+5,
  "maximum_evaluations": 1e+5,
  "atom_style": "atomic"
}
```

Subtask 4 : Slab Relaxation for Surface Energy

Description: Create a slab by removing periodic boundary conditions along the given direction, perform structure relaxation (5×5×5 replication). Report total energy (eV) and saved path.

Input Parameters:

```
{
  "units": "metal",
  "minimiser": "conjugate gradient",
  "energy_tolerance": 1e-10,
  "force_tolerance": 1e-10,
  "maximum_iterations": 1e+5,
  "maximum_evaluations": 1e+5,
  "atom_style": "atomic",
  "direction": [0, 0, 1]
}
```

Subtask 5 : Surface Energy Calculation

Description: Compute surface energy from relaxed bulk and slab structures. Return energy in $\text{eV}/\text{\AA}^2$.

Input Parameters: None

640

Aluminum Melting Simulation

Simulate the melting of Aluminum by first equilibrating the system under NVT conditions, then relaxing it under NPT conditions, followed by a heating stage under NPT where the temperature is gradually increased to induce melting. Use the Embedded Atom Method (EAM) potential and the given simulation parameters. The simulation should replicate the unit cell 5 times in each direction. Report the final density of the system in g/cm^3 .

Input Parameters:

```
{
  "units": "metal",
  "pbc": "true",
  "timestep": "1 femtosecond",
  "nvt steps": 10000,
  "nvt temperature": 300,
  "nvt thermostat": "Nose-Hoover",
  "npt relaxation steps": 10000,
  "npt relaxation temperature": 300,
  "npt relaxation pressure": 0,
  "npt relaxation thermostat": "Nose-Hoover",
  "npt relaxation barostat": "Nose-Hoover",
  "npt heating steps": 20000,
  "npt heating initial temperature": 300,
  "npt heating final temperature": 2000,
  "npt heating pressure": 0,
  "npt_heating_thermostat": "Nose-Hoover",
  "npt_heating_barostat": "Nose-Hoover",
  "materials_project_id": "mp-134",
  "atom_style": "atomic"
}
```

Subtask 1: Retrieve Aluminum Structure

Retrieve structure of Aluminum and convert it to LAMMPS data format. Return the absolute path of the final structure.

Input Parameters:

```
{
  "materials_project_id": "mp-134",
  "atom_style": "atomic"
}
```

Subtask 2: Procure EAM Potential File

Procure the appropriate Embedded Atom Method (EAM) potential file for Aluminum for MD simulations. Return the absolute path to the potential file.

Input Parameters: None

641

Subtask 3: NVT Equilibration

Given paths to LAMMPS-compatible Aluminum structure file and EAM potential file, equilibrate the system under NVT conditions. Replicate the unit cell $5 \times 5 \times 5$. Return the absolute path of the equilibrated structure and final density (g/cm^3).

Input Parameters:

```
{
  "units": "metal",
  "pbc": "true",
  "timestep": "1 femtosecond",
  "nvt steps": 10000,
  "nvt temperature": 300,
  "nvt thermostat": "Nose-Hoover",
  "atom_style": "atomic"
}
```

Subtask 4: NPT Relaxation

Given paths to NVT equilibrated Aluminum structure and EAM potential file, relax the system under NPT conditions. Return the absolute path of the relaxed structure and density (g/cm^3).

Input Parameters:

```
{
  "units": "metal",
  "pbc": "true",
  "timestep": "1 femtosecond",
  "npt relaxation steps": 10000,
  "npt relaxation temperature": 300,
  "npt relaxation pressure": 0,
  "npt relaxation thermostat": "Nose-Hoover",
  "npt relaxation barostat": "Nose-Hoover",
  "atom_style": "atomic"
}
```

Subtask 5: NPT Heating / Melting Simulation

Given paths to NPT equilibrated Aluminum structure and EAM potential file, simulate melting by heating under NPT to gradually increase temperature. Report the final density (g/cm³).

Input Parameters:

```
{
  "units": "metal",
  "pbc": "true",
  "timestep": "1 femtosecond",
  "npt heating steps": 20000,
  "npt heating initial temperature": 300,
  "npt heating final temperature": 2000,
  "npt heating pressure": 0,
  "npt_heating_thermostat": "Nose-Hoover",
  "npt_heating_barostat": "Nose-Hoover",
  "atom_style": "atomic"
}
```

643

Aluminum Quenching Simulation

Simulate the quenching process of Aluminum by cooling its molten structure under NPT conditions, followed by relaxation under NPT and NVT conditions. Use the Embedded Atom Method (EAM) potential and the given simulation parameters. Report the final density of the system in g/cm³.

Input Parameters:

```
{
  "path to molten structure": "/structures/task_30/
    heatedsystem_3000_reax.dat",
  "units": "metal",
  "pbc": "true",
  "timestep": "1 femtosecond",
  "npt quench steps": 20000,
  "npt quench initial temperature": 2000,
  "npt quench final temperature": 300,
  "npt quench pressure": 0,
  "npt quench thermostat": "Nose-Hoover",
  "npt quench barostat": "Nose-Hoover",
  "npt relaxation steps": 10000,
  "npt relaxation temperature": 300,
  "npt relaxation pressure": 0,
  "npt relaxation thermostat": "Nose-Hoover",
  "npt relaxation barostat": "Nose-Hoover",
  "nvt steps": 10000,
  "nvt temperature": 300,
  "nvt thermostat": "Nose-Hoover",
  "materials_project_id": "mp-134",
  "atom_style": "atomic"
}
```

644

Subtask 1: Procure EAM Potential File

Procure the appropriate Embedded Atom Method (EAM) potential file for Aluminum for MD simulations. Return the absolute path to the potential file.

Input Parameters: None

Subtask 2: NPT Quenching

Given paths to LAMMPS-compatible molten Aluminum structure and EAM potential file, simulate the quenching process by cooling the molten structure under NPT conditions. Return the absolute path of the cooled structure and the density in g/cm^3 .

Input Parameters:

```
{
  "path to molten structure": "/structures/task_30/
    heatedsystem_3000_reax.dat",
  "units": "metal",
  "pbc": "true",
  "timestep": "1 femtosecond",
  "npt quench steps": 20000,
  "npt quench initial temperature": 2000,
  "npt quench final temperature": 300,
  "npt quench pressure": 0,
  "npt quench thermostat": "Nose-Hoover",
  "npt quench barostat": "Nose-Hoover",
  "atom_style": "atomic"
}
```

Subtask 3: NPT Relaxation

Given paths to LAMMPS-compatible Aluminum structure and EAM potential file, relax the structure under NPT conditions. Return the absolute path of the relaxed structure and density (g/cm^3).

Input Parameters:

```
{
  "units": "metal",
  "pbc": "true",
  "timestep": "1 femtosecond",
  "npt relaxation steps": 10000,
  "npt relaxation temperature": 300,
  "npt relaxation pressure": 0,
  "npt relaxation thermostat": "Nose-Hoover",
  "npt relaxation barostat": "Nose-Hoover",
  "atom_style": "atomic"
}
```


Subtask 4: NVT Relaxation

Given paths to LAMMPS-compatible Aluminum structure and EAM potential file, relax the structure under NVT conditions. Report the final density (g/cm³).

Input Parameters:

```
{
  "units": "metal",
  "pbc": "true",
  "timestep": "1 femtosecond",
  "nvt_steps": 10000,
  "nvt_temperature": 300,
  "nvt_thermostat": "Nose-Hoover",
  "atom_style": "atomic"
}
```

646

647 F.4 OpenCatalyst

648 F.4.1 Tools

649 The tools of the `OPENCATALYST` environment are described in Table F.7. These tools are highly based
650 on functions from the `Pymatgen` package [42].

Table F.7: OpenCatalyst **tools**. This table describes each tool available in the environment, including its arguments, return values, and potential errors.

Tool
<code>get_structure_from_mp_text</code>
Arguments: <code>mp_id</code> (string): Materials Project identifier string. The unique identifier used by Materials Project to catalog materials, in the format “mp-XXXXXX”.
Return: Returns a CIF content string containing the crystal structure data, including lattice parameters, atomic positions, space group, and symmetry operations.
Raise: <code>ConnectionError</code> : If unable to connect to the Materials Project API due to network issues or server downtime. <code>KeyError</code> : If the specified MP ID is not found in the database, indicating an invalid or non-existent material ID. <code>AuthenticationError</code> : If the API key is invalid or missing, preventing access to the Materials Project database.
<code>enumerate_slabs_text</code>

Tool

Arguments: `bulk_cif` (string): Bulk crystal structure in CIF string format.
`miller_index` (tuple): Miller indices for surface orientation. Defaults to (1,1,1).
`min_slab_size` (float): Minimum slab thickness in Angstroms. Defaults to 12.
`min_vacuum_size` (float): Minimum vacuum layer thickness in Angstroms. Defaults to 5.

Return: Returns a JSON string mapping slab indices to their CIF representations.

Raise: `ValueError`: If CIF string is malformed or parameters are invalid.
`StructureError`: If slab generation fails for the given structure.

choose_slab_text

Arguments: `slabs_json` (string): JSON string mapping slab keys to CIF strings. A JSON-formatted string containing a dictionary where keys are slab identifiers (e.g., “slab_0”, “slab_1”) and values are the corresponding CIF strings.
`index` (int, optional): Index of the slab to select. Defaults to 0. The numerical index of the slab to select from the JSON dictionary.

Return: Returns a CIF string for the selected slab, containing atomic positions, lattice parameters, and other crystallographic information.

Raise: `ValueError`: Raised if the specified slab index is not found in the JSON dictionary.
`JSONDecodeError`: Raised if the `slabs_json` string is not valid JSON.

get_adsorption_sites_text

Arguments: `slab_cif` (string): CIF string of the surface slab structure, including atomic positions, lattice parameters, and surface geometry.

Return: Returns a JSON string containing classified adsorption sites with fractional coordinates.

Raise: `ValueError`: If the CIF string is malformed or doesn’t represent a valid slab.
`StructureError`: If the slab structure cannot be analyzed for adsorption sites due to geometric issues.

choose_adsorption_site_text

Arguments: `adsorption_sites_json` (string): A valid JSON string containing classified adsorption sites, where keys are site types (e.g., “ontop”, “bridge”, “hollow”) and values are lists of fractional coordinates.
`site_type` (string): The type of adsorption site to select (e.g., “ontop”, “bridge”, “hollow”).
`index` (integer, optional): The index of the site to select within the specified site type. Defaults to 0.

Return: Returns a list of three floating-point numbers representing the fractional coordinates [x, y, z] of the selected adsorption site.

Raise: `ValueError`: If the specified `site_type` is not found in the JSON string.
`IndexError`: If the specified index is out of range for the available sites of the specified `site_type`.
`JSONDecodeError`: If the `adsorption_sites_json` string is not valid JSON.

Tool
add_adsorbate_to_slab_text
Arguments: slab_cif (string): CIF string of the surface slab structure. adsorbate_cif (string): CIF or XYZ string of the adsorbate molecule structure. height (float, optional): Height in Angstroms above the surface for adsorbate placement (default is 2.0). site (list of floats or None, optional): Fractional coordinates for adsorbate placement [x, y, z], or None to auto-select site (default is None). Return: Returns a CIF string of the combined surface-adsorbate structure. Raise: ValueError: If CIF strings are malformed or the adsorbate cannot be parsed. StructureError: If adsorbate placement fails due to geometric constraints. RuntimeError: If no adsorption sites are found on the surface.

651 **F.4.2 Scoring Functions**

652 Table F.8 describes the different scoring function used for scoring the OPENCATALYST environment
653 within the *Corral* environment.

Table F.8: OPENCATALYST scoring functions, their arguments, and return values.

Function
resolve_path
Arguments: path_or_str (str): Path or string to resolve. Return: Returns the resolved path as a string.
check_valid_json_file
Arguments: json_path (str): File path to the JSON file. Return: Returns 1.0 if a valid, non-empty JSON file exists at the path; otherwise, 0.0.
check_slabs_json
Arguments: slabs_json (str): JSON string or file path containing slab data. Return: Returns 1.0 if the JSON contains at least one valid slab (parsable CIF); otherwise 0.0.
check_mp_structure
Arguments: path_or_cif (str): CIF string or file path. Return: Returns 1.0 if a valid CIF structure from Materials Project; otherwise, 0.0.
check_slab_structure
Arguments: path_or_cif (str): CIF string or file path. Return: Returns 1.0 if a valid slab structure; otherwise, 0.0.

Function
<code>check_co2_molecule_structure</code> Arguments: <code>path_or_cif</code> (str): CIF string or file path. Return: Returns 1.0 if the structure contains exactly one CO₂ molecule (1 C, 2 O atoms); otherwise, 0.0.
<code>check_adsorption_structure</code> Arguments: <code>slab_elements</code> (list[str]): List of slab elements. <code>adsorbate_elements</code> (list[str]): List of adsorbate elements. Return: Returns a scoring function that scores 1.0 if the structure contains both slab and adsorbate elements; otherwise, it scores 0.0.
<code>check_adsorption_sites</code> Arguments: <code>sites_json_or_path</code> (str): JSON string or file path containing adsorption sites. Return: Returns 1.0 if the adsorption sites JSON contains at least one recognized site type with valid 3D coordinates; otherwise, 0.0.

654 F.4.3 Tasks and Subtasks

655 Here we provide an example of a task and its associated subtask prompts for CO₂ adsorption on
656 Cu₂O. The prompt structure is similar for other compounds.

CO₂ Adsorption on Cu₂O Slab

Create a CO₂ adsorbed structure on a Silicon/Cu₂O slab. Submit the path to the final combined structure CIF file.

Input Parameters:

```
{
  "mp_id": "mp-149",
  "co2_mp_id": "mp-644607",
  "miller_index": [1, 1, 1],
  "min_slab_size": 12,
  "min_vacuum_size": 5,
  "slab_index": 0,
  "site_type": "ontop",
  "site_index": 0,
  "height": 2.0
}
```

657

Subtask 1: Retrieve Bulk Structure

Retrieve structure of Cu_2O from Materials Project and save it as a CIF file. Submit the path to the CIF file.

Input Parameters:

```
{  
  "mp_id": "mp-361"  
}
```

Subtask 2: Enumerate Slabs

Enumerate possible slabs from the bulk Cu_2O structure and save the result as a JSON file. Submit the path to the JSON file.

Input Parameters:

```
{  
  "miller_index": [1, 1, 1],  
  "min_slab_size": 12,  
  "min_vacuum_size": 5  
}
```

Subtask 3: Select Slab

Choose one slab from the enumerated slabs (by index) and save it as a CIF file. Submit the path to the CIF file.

Input Parameters:

```
{  
  "index": 0  
}
```

Subtask 4: Retrieve CO_2 Molecule

Retrieve CO_2 molecule structure from Materials Project and save it as a CIF file. Submit the path to the CIF file.

Input Parameters:

```
{  
  "mp_id": "mp-644607"  
}
```

Subtask 5: Determine Adsorption Sites

Determine possible adsorption sites on the chosen slab and save the results as a JSON file. Submit the path to the JSON file.

Input Parameters: None

Subtask 6: Select Adsorption Site

Choose one adsorption site (preferably ontop site) from the identified sites and save the coordinates to a file. Submit the path to the file.

Input Parameters:

```
{
  "site_type": "ontop",
  "index": 0
}
```

Subtask 7: Place CO₂ on Slab

Place the CO₂ molecule on the chosen slab at the specified adsorption site with a height of approximately 2.0 Å and save the combined structure as a CIF file. Submit the path to the CIF file.

Scoring Parameters:

```
{
  "slab_elements": ["Cu", "O"],
  "adsorbate_elements": ["C", "O"]
}
```

659

660 **F.5 ML**

661 **F.5.1 Tools**

662 Table F.9 shows the different tools available for the agents when working in the ML environment.

Table F.9: ML tools, their arguments, and return values.

Tool
<code>get_bulk_polymorphs_data</code>
Arguments: <code>composition</code> (string): Chemical formula specifying the composition for which polymorphs should be retrieved. Should follow standard chemical notation (e.g., TiO ₂ , Al ₂ O ₃).
Return: Returns a JSON string containing comprehensive polymorph data, sorted by energy above hull. The data includes properties like Materials Project ID, CIF structure, energy above hull, formation energy per atom, band gap, density, volume, number of sites, space group, and stability information.
Raise: <code>KeyError</code> : Raised when the specified composition is not found in the database. <code>ValueError</code> : Raised when the Materials Project API key is not available. <code>ConnectionError</code> : Raised when there is a failure to connect to the Materials Project API.
<code>get_bulk_polymorphs_data_to_file</code>

Tool

Arguments: `composition` (string): Chemical formula specifying the composition for which polymorphs should be retrieved. Example: TiO_2 (titanium dioxide), " Al_2O_3 " (aluminum oxide).

`save_path` (string or None): Path where JSON data will be saved. Must include the filename with `.json` extension. If None, an error is raised.

Return: Returns the file path where polymorph data was saved in JSON format.

Raise: `ValueError`: If `save_path` is None or API key is missing.

`IOError`: If unable to write to the specified file path (e.g., directory does not exist or is not writable).

batch_retrieve_polymorphs

Arguments: `compositions` (list of strings): List of chemical compositions to retrieve polymorphs for.

`max_energy_above_hull` (float, optional): Maximum energy above hull threshold in eV/atom. Defaults to 0.5.

`max_per_composition` (int, optional): Maximum number of polymorphs per composition. Defaults to 10.

`save_directory` (string, optional): Directory path for saving individual composition files. Defaults to "polymorph_data".

`work_dir` (string, optional): Working directory for the tool (hidden argument).

Return: Returns a JSON string containing batch retrieval results and statistics, including successful and failed compositions, total number of polymorphs, and file paths for each composition.

Raise: Raises `ValueError` if parameters are invalid (negative energy, zero `max_per_composition`). Raises `IOError` if unable to create the save directory or write files.

sort_and_get_first_from_json

Arguments: `polymorph_data_json` (string): JSON string containing the data to be sorted, structured as a list of dictionaries with numerical properties for sorting.

`sort_key` (string): Property name to sort the data by (e.g., "energy_above_hull", "band_gap", "density").

`return_key` (string): Property name to return from the first element after sorting (e.g., "material_id", "cif").

Return: Returns the value of the specified `return_key` from the first element after sorting.

Raise: `JSONDecodeError`: If the `polymorph_data_json` is not valid JSON.

`KeyError`: If `sort_key` or `return_key` are not found in the data.

`IndexError`: If the JSON data is empty or contains no elements.

select_polymorphs_with_strategy

Tool

Arguments: `polymorphs_data` (string): JSON string or file path containing polymorph data.
 `selection_strategy` (string, optional): “diverse_energy”, “most_stable”, or “diverse_structure” to determine selection method. Defaults to “diverse_energy”.
 `max_polymorphs` (int, optional): Maximum number of polymorphs to select. Defaults to 5.
 `energy_threshold` (float, optional): Maximum energy above hull in eV/atom. Defaults to 0.5.
 `is_path` (bool, optional): Whether `polymorphs_data` is a file path. Defaults to False.

Return: Returns a JSON string containing selected polymorphs based on the specified strategy.

Raises: `ValueError`: If invalid selection strategy is specified.
 `FileNotFoundError`: If `is_path=True` but file doesn’t exist.
 `JSONDecodeError`: If `polymorphs_data` contains invalid JSON.

consolidate_polymorph_datasets

Arguments: `composition_files` (dict): Dictionary mapping compositions to their JSON file paths.
 `output_path` (string, optional): Path for the consolidated dataset (default: “consolidated_polymorphs.json”).
 `work_dir` (string, optional): Working directory for intermediate steps (hidden argument).

Return: Returns a JSON-formatted string with consolidation results and detailed statistics (e.g., total polymorphs, number of compositions included, etc.).

Raise: `FileNotFoundError`: If one or more input files cannot be found.
 `JSONDecodeError`: If input files contain invalid JSON.
 `IOError`: If unable to write to the output path.

select_polymorphs_with_strategy_to_file

Tool

Arguments: `polymorphs_data` (string): JSON string or file path containing polymorph data.
`save_path` (string): File path where selected polymorphs will be saved in JSON format.
`selection_strategy` (string): selection strategy (“diverse_energy”, “most_stable”, “diverse_structure”). Default is `diverse_energy`.
`max_polymorphs` (int): Maximum number of polymorphs to select. Default is 5.
`energy_threshold` (float): Maximum energy above hull in eV/atom. Default is 0.5.
`is_path` (bool): Whether `polymorphs_data` is a file path. Default is `False`.
`work_dir` (string | None): Directory to work in (optional).
Return: Returns the file path where the selected polymorphs were saved.
Raise: `ValueError`: If an invalid selection strategy is specified.
`FileNotFoundError`: If `is_path=True` but the input file doesn’t exist.
`IOError`: If unable to write to the `save_path` location.

`filter_json_with_strategy`

Arguments: `input_json_path` (string): Path to the input JSON file to be filtered.
`output_json_path` (string): Path where filtered JSON data will be saved.
`custom_code` (string | None): Python code string defining the filtering logic.
`work_dir` (string | None): Directory to work in, optional.
Return: Returns a JSON string with filtering results and comprehensive statistics, including original and filtered data counts, output file path, and reduction percentage.
Raise: `FileNotFoundError`: Raised when the input JSON file doesn’t exist.
`JSONDecodeError`: Raised when the input file contains invalid JSON.
`SyntaxError`: Raised when the custom filtering code contains syntax errors.
`RuntimeError`: Raised when the custom filtering code fails during execution.

`prepare_tabular_dataset`

Tool

Arguments: `polymorphs_json_path` (string): Path to consolidated polymorphs JSON file containing material properties and crystal structures.
`output_path` (string): Base directory and filename prefix for saving the dataset files.
`target_property` (string, optional): Property to predict (defaults to “formation_energy_per_atom”).
`feature_engineering` (string): feature engineering strategy (“basic”, “advanced”, “custom”).
`test_split` (float, optional): Fraction of data for test set (defaults to 0.2).
`normalize` (bool, optional): Whether to normalize features (defaults to True).
`work_dir` (string, optional): Working directory for intermediate files.

Return: Returns a JSON string with dataset preparation results, including file paths for training and test data, and other metadata.

Raise: `FileNotFoundError` if the polymorphs JSON file is not found.
`KeyError` if the target property is not found in the data.
`ValueError` if feature engineering fails or the data format is invalid.

`evaluate_xgboost_model`

Arguments: `model_path` (string): Path to the saved XGBoost model file.
`test_data_path` (string): Path to the test data CSV file with the same structure as the training data.
`target_column` (string, default = “formation_energy_per_atom”): The column name in the test data representing the target variable for evaluation.
`detailed_analysis` (bool, default = True): Whether to include detailed analysis and feature importance in the results.

Return: Returns a JSON string containing evaluation success status and metrics, including MAE, RMSE, R2, MAPE, prediction ranges, error analysis, and feature importance (if detailed analysis is enabled).

Raise: `FileNotFoundError`: Raised when the model file or test data file doesn’t exist.
`KeyError`: Raised when the target column is not found in the test data.
`ValueError`: Raised when the model and data are incompatible or contain invalid values.

`train_xgboost_model`

Tool
Arguments: <code>train_data_path</code> (string): Path to the training data CSV file containing features and target column. <code>test_data_path</code> (string): Path to the test data CSV file, with the same structure as the training data. <code>model_save_path</code> (string): Path to save the trained model file in .pkl format. <code>target_column</code> (string, default="formation_energy_per_atom"): Name of the target column for prediction in the dataset. <code>hyperparameters</code> (dict None, optional): Dictionary of XGBoost hyperparameters to override default values. <code>work_dir</code> (string None, optional): Working directory path, if provided. Return: Returns a JSON-formatted string with comprehensive training results, including model performance metrics (MAE, RMSE, R2), feature importance, and the model path. Raises: <code>FileNotFoundError</code>: Raised when the specified CSV files do not exist. <code>KeyError</code>: Raised when the target column does not exist in the dataset. <code>ValueError</code>: Raised when the data contains invalid values or formatting issues.
perform_cross_validation
Arguments: <code>train_data_path</code> (string): Path to the training data CSV file. <code>target_column</code> (string): Name of the target column for prediction (default: "formation_energy_per_atom"). <code>cv_folds</code> (int): Number of cross-validation folds (default: 5). <code>hyperparameters</code> (dict or None): XGBoost hyperparameters for cross-validation (default: None). Return: Returns a JSON string containing comprehensive cross-validation results and statistical analysis, including R2, MAE, and other performance metrics. Raise: <code>FileNotFoundError</code>: Raised if the training data file does not exist or is not found. <code>KeyError</code>: Raised if the target column is not found in the dataset. <code>ValueError</code>: Raised if the number of cross-validation folds is invalid or the data contains invalid values.

663 F.5.2 Scoring Functions

664 Table F.10 show and describe the scoring functions used in the ML environment.

Table F.10: **Scoring Functions for the ML Environment.** The table describes each scoring function, its arguments, return values, and potential behaviors.

Function
<code>check_mp_structure</code> Arguments: <code>path_or_cif</code> (string): Path to CIF file or CIF string. Return: Returns 1.0 if the input is a valid CIF string or file representing a non-empty structure (via <code>pymatgen.Structure</code>), otherwise 0.0.

Function
<code>compare_with_ground_truth</code>
Arguments: <code>generated_path</code> (string): Path to generated JSON file. <code>ground_truth_path</code> (string): Path to reference JSON file. <code>comparison_mode</code> (string, optional): “strict“, “keys“, “numerical“, or “subset“. <code>tolerance</code> (float, optional): Numerical tolerance. Return: Returns 1.0 if the generated file matches the ground truth according to the chosen mode, otherwise 0.0. Modes control the exactness of comparison (structure, keys, numerical values, subset).
<code>ml_pipeline_score</code>
Arguments: <code>model_path</code> (string): Path to trained ML model file. Return: Returns a float score (0.0-1.0) evaluating the completeness and performance of an ML pipeline, including model loading, evaluation metrics, and cross-validation evidence.
<code>polymorph_retrieval_success</code>
Arguments: <code>retrieval_results_path</code> (string): Path to JSON retrieval results. Return: Returns a float score (0.0-1.0) based on polymorph retrieval success rate, total number retrieved, and distribution per composition.
<code>score_polymorph_dataset</code>
Arguments: <code>consolidated_json_path</code> (string): Path to consolidated polymorph dataset JSON. Return: Returns 1.0 if at least one composition has multiple polymorphs, otherwise 0.0.
<code>ml_dataset_preparation_quality_binary</code>
Arguments: <code>ml_metadata_path</code> (string): Path to ML dataset metadata JSON. Return: Returns 1 if dataset preparation meets all criteria: existing train/test files, sufficient sample sizes, and adequate feature count; otherwise returns 0.
<code>model_training_success_binary</code>
Arguments: <code>model_path</code> (string): Path to trained model file. Return: Returns 1 if model can be loaded, has a <code>predict</code> method, and meets performance thresholds ($R^2 \geq 0.5$, $MAE \leq 0.5$), otherwise 0.
<code>model_evaluation_completeness_binary</code>
Arguments: <code>evaluation_results_path</code> (string): Path to model evaluation results JSON. Return: Returns 1 if evaluation results are complete: includes required metrics (MAE, RMSE, R^2), $R^2 \geq 0.7$, cross-validation stats, feature importance, and detailed analysis; otherwise returns 0.

666 Here we provide an example of a task and its associated subtask prompts for oxide. The prompt
667 structure is similar for sulphide and nitride compounds.

Oxide Polymorph Dataset and XGBoost Model

Generate a comprehensive dataset of oxide polymorphs from Materials Project and train an XGBoost model to predict formation energies. Evaluate the trained XGBoost model using test set and cross-validation metrics. Save the results as a JSON file with keys `test_set_evaluation` and `cross_validation_results`. The `test_set_evaluation` dictionary must contain `mae`, `rmse`, `r2`, and `feature_importance`. The `cross_validation_results` dictionary must contain `r2_mean` and `r2_std`.

Input Parameters:

```
{
  "max_energy_above_hull": 0.3,
  "max_per_composition": 3,
  "max_compositions": 40,
  "target_property": "formation_energy_per_atom",
  "feature_engineering": "advanced",
  "test_split": 0.2,
  "normalize": true,
  "n_estimators": 200,
  "max_depth": 8,
  "learning_rate": 0.1,
  "subsample": 0.8,
  "colsample_bytree": 0.8,
  "reg_alpha": 0.1,
  "reg_lambda": 0.1,
  "detailed_analysis": true,
  "cv_folds": 5
}
```

Subtask 1: Generate Oxide Polymorph Dataset

Generate a diverse list of oxide compositions for dataset creation. Retrieve polymorphs for all oxide compositions from Materials Project database. For each composition, collect multiple polymorphs including both stable and metastable structures. Consolidate all individual polymorph files into a single comprehensive dataset. Each entry should have `source_composition`.

Input Parameters:

```
{
  "max_energy_above_hull": 0.3,
  "max_per_composition": 3,
  "max_compositions": 40
}
```

Subtask 2: Prepare ML Dataset

Transform the consolidated dataset into ML-ready format with engineered features and proper train/test splits. Create a metadata JSON file with features, train_path, and test_path.

Input Parameters:

```
{
  "target_property": "formation_energy_per_atom",
  "feature_engineering": "advanced",
  "test_split": 0.2,
  "normalize": true,
  "output_path": "oxide_ml_dataset"
}
```

Subtask 3: Train XGBoost Model

Train an XGBoost regression model to predict formation energies of oxide polymorphs. Use the prepared dataset with optimized hyperparameters for oxide materials. Focus on achieving good generalization performance across different oxide families and structural types.

Input Parameters:

```
{
  "target_column": "formation_energy_per_atom",
  "hyperparameters": {
    "n_estimators": 200,
    "max_depth": 8,
    "learning_rate": 0.1,
    "subsample": 0.8,
    "colsample_bytree": 0.8,
    "reg_alpha": 0.1,
    "reg_lambda": 0.1
  }
}
```

Subtask 4: Evaluate Model

Evaluate the trained XGBoost model using test set and cross-validation metrics. Save the results as a JSON file with keys test_set_evaluation and cross_validation_results. The test_set_evaluation dictionary must contain mae, rmse, r2, and feature_importance. The cross_validation_results dictionary must contain r2_mean and r2_std.

Input Parameters:

```
{
  "target_column": "formation_energy_per_atom",
  "detailed_analysis": true,
  "cv_folds": 5
}
```

670 G Pass^k

671 Beyond the standard pass@5 metric, we also computed the pass⁵ score (more details about the
 672 metrics are presented in Appendix B). The results, presented in Figure G.2 for the four environments,
 673 two agents, and two LLMs, reveal trends highly consistent with those of pass@5 (see Figure 2).
 674 Specifically, SPECTRA and MD perform notably poorly, while OPENCATALYST achieves near-perfect
 675 scores. A key observation, however, is the pronounced decline in performance for the ML environment
 676 under the pass⁵ metric. Its scores decreased substantially from the pass@5 assessment, falling to a
 677 level comparable with those of the MD environment. This, we hypothesize, may be because of the
 678 distance between tasks being comparable to those of the SPECTRA environment. As for pass@5, no
 679 trends are observed among the different agents and LLMs that were evaluated.




Figure G.2: **pass⁵ scores across evaluated environments, agents, and LLMs.** Performance trends largely align with those of pass@5, with OPENCATALYST achieving near-perfect scores and SPECTRA and MD performing poorly. A pronounced deviation is observed for the ML environment, where a substantial performance decline under the pass⁵ metric suggests agents struggle to solve its tasks consistently. This decline may be attributed to the increased inter-task distance, which appears comparable to that of the SPECTRA environment. No significant trends were identified across different agents or LLMs.

H Number of Tool Calls

The number of tool calls executed by the agents constitutes another metric captured in *Corral*. While straightforward to quantify, this metric offers valuable insights into an agent’s performance and behavioral patterns. As illustrated in Figure H.3, the number of tool calls varies significantly across environments. An inverse relationship is observed between the frequency of tool use and both overall performance and the level of required domain expertise. Specifically, the SPECTRA and MD environments exhibit the highest volume of tool calls. In contrast, the number of tool calls is considerably lower and comparable between the OPENCATALYST and ML environments. Furthermore, the tool-calling agent consistently employs a greater number of tool calls than its ReAct counterpart. This discrepancy is anticipated, as the tool-calling agent’s strategy depends more heavily on tool sampling and less on explicit reasoning chains.




Figure H.3: **Total tool calls per environment, sorted by required domain expertise.** Environments demanding less expertise (OPENCATALYST, ML) exhibit fewer tool calls, while those requiring more (SPECTRA, MD) prompt significantly higher usage. The tool-calling agent consistently employs more calls than the ReAct agent across all environments.

H.1 Performance and Number of Tools

The observed trend suggests a potential inverse relationship between the number of tool calls and pass@5 scores, a comparison visually supported by Figure H.4. This figure indicates that high performance on the OPENCATALYST and ML environments is associated with a low number of tool calls. Conversely, lower scores on the MD and, more markedly, on the SPECTRA environments coincide with a substantially higher frequency of tool calls. One might hypothesize that this trend is driven by an inherent requirement for more steps or a greater number of available tools in certain environments. However, this explanation is not supported by the data. The number of steps required for SPECTRA, MD, and OPENCATALYST is comparable, particularly as both MD and OPENCATALYST involve numerous file management operations absent in SPECTRA. Similarly, the possibility that a larger toolkit acts as a distractor is countered by the fact that the number of tools available in SPECTRA is comparable to that in ML, as detailed in Table 1.




Figure H.4: Illustration of the inverse trend between pass@5 scores and tool call frequency. The high scores in OPENCATALYST and ML correspond to low tool usage, while the low scores in MD and SPECTRA correspond to high, inefficient tool usage. This divergence occurs despite comparable procedural steps and toolset sizes, countering hypotheses that attribute performance solely to environmental complexity.

703 I Number of Completion Tokens

704 Although the number of tool calls demonstrates a strong correlation with model performance, the
 705 consumption of completion tokens does not exhibit a similar trend, as illustrated in Figure I.5. The
 706 count of completion tokens remains largely consistent across the majority of environments and
 707 agents. A prominent exception is the MD environment, where the ReAct agent implemented with
 708 Claude-3.5 Sonnet generates a substantially higher volume of tokens. Apart from this outlier, the
 709 ML environment is the only one that consistently results in a lower token count relative to the others.
 710 Regarding agent performance, it is noteworthy that the ReAct agent does not systematically produce
 711 more tokens despite its requirement to articulate its reasoning through explicit “thoughts”.




Figure I.5: **Number of completion tokens across environments sorted by the domain expertise required to solve them.** The figure illustrates that the number of tokens is similar across environments and agents. Only the ML environment seems to yield a lower number of completion tokens.

J Tool verbosity ablation results

Based on the initial results obtained across the environments, we investigated whether the verbosity level employed in the tool descriptions provided to the agent could significantly influence agent performance. As shown in Table J.11, no clear trend was observed across the different verbosity levels, with the performance trends across environments remaining largely consistent. The most notable performance difference occurred in the MD environment at the COMPREHENSIVE verbosity level when compared to the other two description levels. This difference, however, was less pronounced for the chained tasks.

Table J.11: **Ablation Study on Tool Description Verbosity.** This table presents the results of varying the level of detail in tool descriptions provided to the agents (brief, workflow, and comprehensive). The text in bold shows the best scores across environments. In the simpler OPENCATALYST and ML environments, performance is consistently high regardless of verbosity. However, in the more complex MD and SPECTRA domains, providing more context through workflow or comprehensive descriptions generally improves scores, indicating that richer information about the action space is crucial for the agent’s policy to succeed in challenging tasks.

Agent	Verbosity	MD		ML		OPENCATALYST		SPECTRA	
		single	chained	single	chained	single	chained	single	chained
Claude 3.5 Sonnet									
ReAct	brief	0.22	0.5	0.89	0.96	1.0	1.0	0.19	0.47
	workflow	0.21	0.51	0.56	0.91	1.0	1.0	0.18	0.43
	comprehensive	0.57	0.69	0.97	0.97	1.0	1.0	0.24	0.44
Tool-Calling	brief	0.52	0.78	0.99	0.91	1.0	0.66	0.05	0.49
	workflow	0.81	0.74	1.0	0.99	1.0	1.0	0.1	0.46
	comprehensive	0.69	0.65	0.89	1.0	1.0	1.0	0.17	0.3
GPT-4o									
ReAct	brief	0.2	0.72	0.33	0.99	0.89	1.0	0.05	0.44
	workflow	0.51	0.71	0.89	1.0	1.0	1.0	0.0	0.38
	comprehensive	0.42	0.72	0.97	1.0	1.0	1.0	0.0	0.36
Tool-Calling	brief	0.2	0.69	0.33	0.77	1.0	0.75	0.05	0.41
	workflow	0.07	0.77	1.0	0.97	1.0	1.0	0.0	0.39
	comprehensive	0.19	0.73	1.0	1.0	1.0	1.0	0.0	0.41

K Embeddings distance ablation

For studying the consistently poor performance of agents in the SPECTRA environment, we decided to analyze the cosine distance of tasks and tools from that environment. We used the embeddings obtained through the API of the closed-source model text-embedding-3-large.

For obtaining the distances reported in Table 1, each task and tool description is passed through the model to generate its respective embedding. These embeddings are then used to calculate pairwise distances between tools, tasks, and the combinations of tools and tasks. The distance metric used for comparison is the cosine distance, which is calculated as follows:

$$D_{\text{cosine}}(A, B) = 1 - \frac{A \cdot B}{\|A\|_2 \|B\|_2}$$

where:

- A and B are the vectors (embeddings) representing the two entities (tasks or tools),
- $\|A\|_2$ and $\|B\|_2$ are the L2 norms (magnitudes) of the vectors,
- $A \cdot B$ is the dot product between the vectors.

732 L Overall Results

Table L.12: Performance comparison of agents across SPECTRA, ML and OPENCATALYST tasks with varying tool verbosity.

Task/ Env	Task Level	Tool Verbosity	Agent	Model	Scores			Tool Exec time (s)	Total tokens	Tool calls	
					Overall	pass@5	pass^5			Total	Failed
SPECTRA	Task	brief	Tool-calling	Claude 3.5	0.050	0.050	0.050	2517.794	1999844	316	0
			Tool-calling	GPT-4o	0.020	0.046	0.001	4954.448	1066725	644	0
			ReAct	Claude 3.5	0.140	0.192	0.101	2761.069	1075978	487	0
		comprehensive	ReAct	GPT-4o	0.020	0.046	0.001	4669.816	2982094	855	0
			Tool-calling	Claude 3.5	0.120	0.167	0.100	3321.458	9600785	681	0
			Tool-calling	GPT-4o	0	0.000	0.000	3986.564	7067824	1350	0
		workflow	ReAct	Claude 3.5	0.150	0.238	0.068	4672.958	11309758	890	0
			ReAct	GPT-4o	0	0.000	0.000	3372.496	8579030	575	0
			Tool-calling	Claude 3.5	0.100	0.100	0.100	2909.819	6487221	680	0
			Tool-calling	GPT-4o	0	0.000	0.000	5425.582	4845761	1284	0
			ReAct	Claude 3.5	0.140	0.183	0.104	3491.533	7418732	841	0
			ReAct	GPT-4o	0	0.000	0.000	3555.800	5418950	632	1
ML	Task	brief	Tool-calling	Claude 3.5	0.447	0.492	0.407	14603.170	3052471	1591	0
			Tool-calling	GPT-4o	0.343	0.415	0.286	25426.708	3711126	1797	0
			ReAct	Claude 3.5	0.416	0.467	0.372	12952.609	1845878	1221	0
		comprehensive	ReAct	GPT-4o	0.366	0.436	0.303	13407.503	2414056	1510	0
			Tool-calling	Claude 3.5	0.281	0.304	0.257	2512.472	7929146	1753	20
			Tool-calling	GPT-4o	0.344	0.407	0.294	13817.494	7819026	1857	0
		workflow	ReAct	Claude 3.5	0.410	0.436	0.377	11667.673	7265056	1480	48
			ReAct	GPT-4o	0.332	0.365	0.290	6221.911	6186909	1329	10
			Tool-calling	Claude 3.5	0.417	0.460	0.370	2625.681	7404750	1740	16
			Tool-calling	GPT-4o	0.333	0.388	0.299	14822.657	5269491	1658	0
			ReAct	Claude 3.5	0.388	0.429	0.351	11366.324	5174901	1406	90
			ReAct	GPT-4o	0.331	0.384	0.276	12005.337	7266929	1282	38
	Subtask	brief	Tool-calling	Claude 3.5	0.667	0.993	0.161	1220.296	1845649	147	9
			Tool-calling	GPT-4o	0.267	0.333	0.109	1266.829	1340868	137	6
			ReAct	Claude 3.5	0.533	0.887	0.135	1034.641	933964	116	14
		comprehensive	ReAct	GPT-4o	0.267	0.333	0.109	3337.430	4463267	128	2
			Tool-calling	Claude 3.5	0.600	0.887	0.359	1859.059	5304033	159	7
			Tool-calling	GPT-4o	0.733	0.996	0.244	1176.306	1968853	95	0
		workflow	ReAct	Claude 3.5	0.733	0.974	0.446	1093.865	2211623	96	13
			ReAct	GPT-4o	0.533	0.967	0.055	1097.562	1848462	86	0
			Tool-calling	Claude 3.5	1	1.000	1.000	1425.824	2997344	132	0
			Tool-calling	GPT-4o	0.933	1.000	0.776	1189.020	1600247	103	0
			ReAct	Claude 3.5	0.400	0.557	0.333	1048.410	1731643	99	14
			ReAct	GPT-4o	0.733	0.891	0.667	823.038	1375619	94	4
OPENCATALYST	Task	brief	Tool-calling	Claude 3.5	0.783	0.909	0.589	1024.336	1444605.667	94.333	8.333
			Tool-calling	GPT-4o	0.417	0.772	0.072	2919.916	1283053.667	126.667	12.000
			ReAct	Claude 3.5	0.767	0.959	0.563	628.312	352240.333	76.667	6.333
		comprehensive	ReAct	GPT-4o	0.950	0.994	0.918	572.154	218174.667	49.667	0.000
			Tool-calling	Claude 3.5	0.950	1.000	0.832	602.452	697221.667	65.333	0.000
			Tool-calling	GPT-4o	0.950	0.999	0.867	455.008	511905.000	56.000	0.000
		workflow	ReAct	Claude 3.5	0.833	0.966	0.666	604.539	525621.000	62.333	7.000
			ReAct	GPT-4o	0.917	0.999	0.755	687.699	408186.333	49.333	1.667
			Tool-calling	Claude 3.5	0.933	0.993	0.861	741.300	730223.000	65.000	0.000
			Tool-calling	GPT-4o	0.883	0.966	0.834	414.070	364710.333	54.667	0.000
			ReAct	Claude 3.5	0.833	0.915	0.707	670.169	473178.000	68.333	8.000
			ReAct	GPT-4o	0.967	0.999	0.923	611.969	321803.333	48.333	2.333
	Subtask	brief	Tool-calling	Claude 3.5	0.933	1.000	0.776	2037.189	2711823	188	32
			Tool-calling	GPT-4o	1	1.000	1.000	2130.001	1953905	144	24
			ReAct	Claude 3.5	0.733	0.891	0.667	2290.416	1538987	114	2
		comprehensive	ReAct	GPT-4o	1	1.000	1.000	1702.736	1786819	125	0
			Tool-calling	Claude 3.5	0.933	1.000	0.776	1763.970	3269313	157	37
			Tool-calling	GPT-4o	1	1.000	1.000	2054.796	2314210	120	0
		workflow	ReAct	Claude 3.5	1	1.000	1.000	1705.277	2586206	119	0
			ReAct	GPT-4o	0.933	1.000	0.776	2642.857	2563808	119	0
			Tool-calling	Claude 3.5	1	1.000	1.000	1744.650	2646049	154	26
			Tool-calling	GPT-4o	1	1.000	1.000	2102.629	2357448	153	30
			ReAct	Claude 3.5	1	1.000	1.000	1762.177	2332113	125	0
			ReAct	GPT-4o	0.933	1.000	0.776	2347.130	2064826	160	0
	Subtask	brief	Tool-calling	Claude 3.5	0.552	0.655	0.449	1292.663	793577	111	23
			Tool-calling	GPT-4o	0.714	0.746	0.682	1206.510	851828	140	39
			ReAct	Claude 3.5	0.971	0.999	0.924	810.621	455240	113	10
		comprehensive	ReAct	GPT-4o	0.943	0.996	0.877	1254.770	574773	103	8
			Tool-calling	Claude 3.5	1.000	1.000	1.000	1350.958	981333	126	10
			Tool-calling	GPT-4o	1.000	1.000	1.000	1146.016	692674	109	4
		workflow	ReAct	Claude 3.5	1.000	1.000	1.000	1109.015	806821	116	1
			ReAct	GPT-4o	1.000	1.000	1.000	1203.005	627433	107	0
			Tool-calling	Claude 3.5	0.895	0.999	0.668	1371.961	888996	121	11
			Tool-calling	GPT-4o	1.000	1.000	1.000	1135.358	617183	109	4
			ReAct	Claude 3.5	1.000	1.000	1.000	1086.880	699299	115	0
			ReAct	GPT-4o	1.000	1.000	1.000	1066.127	545052	98	1

Table L.13: Performance comparison of agents across MD tasks with varying tool verbosity

Task/ Env	Task Level	Tool Verbosity	Agent	Model	Scores			Tool Exec time (s)	Total tokens	Tool calls	
					Overall	pass@5	pass [^] 5			Total	Failed
Surface Energy	Task	brief	Tool-calling	Claude 3.5	0.20	0.57	0.00272	3994.60	2372658	243	46
			Tool-calling	GPT-4o	0.00	0.00	0.00	2668.31	2761722	334	53
			ReAct	Claude 3.5	0.40	0.65	0.25	1150.81	740500	132	10
		comprehensive	ReAct	GPT-4o	0.00	0.00	0.00	1937.86	1711175	348	92
			Tool-calling	Claude 3.5	0.30	0.42	0.25	3436.00	4010181	284	52
			Tool-calling	GPT-4o	0.10	0.23	0.0025	2576.81	2756626	376	55
			ReAct	Claude 3.5	0.60	0.75	0.35	2633.79	2173989	204	0
			ReAct	GPT-4o	0.15	0.25	0.02	1872.54	2686373	254	46
		workflow	Tool-calling	Claude 3.5	0.65	0.75	0.41	3143.25	1669994	234	3
			Tool-calling	GPT-4o	0.00	0.00	0.00	2579.06	2608477	346	67
			ReAct	Claude 3.5	0.00	0.00	0.00	2796.24	1455136	91	4
		Subtask	ReAct	GPT-4o	0.10	0.34	0.0001	1732.40	2116582	309	69
			Tool-calling	Claude 3.5	0.86	0.95	0.70	3675.40	2392176	412	21
			Tool-calling	GPT-4o	0.53	0.66	0.42	16050.29	8819312	817	114
			ReAct	Claude 3.5	0.33	0.48	0.21	2707.52	2015689	320	33
			ReAct	GPT-4o	0.55	0.66	0.47	4862.63	3533316	697	131
			Tool-calling	Claude 3.5	0.86	0.95	0.69	3174.83	2310994	327	7
			Tool-calling	GPT-4o	0.56	0.71	0.42	16846.90	7371726	871	155
			ReAct	Claude 3.5	0.61	0.72	0.50	3100.14	2213461	353	26
			ReAct	GPT-4o	0.54	0.70	0.40	5057.13	5066130	702	143
		workflow	Tool-calling	Claude 3.5	0.78	0.94	0.54	3439.93	2623439	394	22
			Tool-calling	GPT-4o	0.55	0.73	0.45	7761.7	7628950	883	154
			ReAct	Claude 3.5	0.60	0.68	0.52	3228.89	2045466	367	27
			ReAct	GPT-4o	0.56	0.68	0.46	4750.14	4227465	682	121
Melting	Task	brief	Tool-calling	Claude 3.5	0.13	0.45	0.0002	2012.93	1843125	138	44
			Tool-calling	GPT-4o	0.00	0.00	0.00	2015.09	1525141	241	57
			ReAct	Claude 3.5	0.0	0.0	0.0	265.25	35472	0	0
		comprehensive	ReAct	GPT-4o	0.00	0.00	0.00	4501.84	2364025	208	35
			Tool-calling	Claude 3.5	0.33	0.64	0.03	1975.85	2215060	170	34
			Tool-calling	GPT-4o	0.20	0.33	0.026	1171.90	1532063	184	7
			ReAct	Claude 3.5	0.67	0.97	0.36	2479.20	1956353	191	24
			ReAct	GPT-4o	0.00	0.00	0.00	1522.47	1554510	131	8
		workflow	Tool-calling	Claude 3.5	0.33	0.78	0.03	1962.91	1398117	139	13
			Tool-calling	GPT-4o	0.00	0.00	0.00	1132.12	902589	169	8
			ReAct	Claude 3.5	0.47	0.64	0.34	1387.85	622430	80	2
		Subtask	ReAct	GPT-4o	0.07	0.22	0.0001	971.58	765381	132	2
			Tool-calling	Claude 3.5	0.55	0.67	0.44	3243.99	3243724	364	69
			Tool-calling	GPT-4o	0.63	0.78	0.45	21719.57	19734289	583	21
			ReAct	Claude 3.5	0.41	0.64	0.22	3795.09	2994944	380	62
			ReAct	GPT-4o	0.53	0.68	0.41	3598.93	2343615	337	3
			Tool-calling	Claude 3.5	0.51	0.57	0.47	3160.90	4366885	348	84
			Tool-calling	GPT-4o	0.57	0.80	0.41	27788.65	7650030	561	44
			ReAct	Claude 3.5	0.80	0.95	0.69	5811.09	5376726	495	76
			ReAct	GPT-4o	0.53	0.64	0.42	3352.13	3057493	313	3
		workflow	Tool-calling	Claude 3.5	0.52	0.74	0.35	3077.21	3278175	348	58
			Tool-calling	GPT-4o	0.64	0.83	0.44	29058.55	4814488	467	18
			ReAct	Claude 3.5	0.48	0.61	0.40	4047.19	2874521	405	64
			ReAct	GPT-4o	0.60	0.69	0.51	3486.44	1861511	315	1
Quenching	Task	brief	Tool-calling	Claude 3.5	0.40	0.56	0.33	2715.42	4212990	204	58
			Tool-calling	GPT-4o	0.27	0.61	0.007	1633.41	2357544	195	27
			ReAct	Claude 3.5	0.0	0.0	0.0	3276.82	650701	30.0	0.0
		comprehensive	ReAct	GPT-4o	0.27	0.61	0.007	3045.26	5253679	198	50
			Tool-calling	Claude 3.5	0.87	0.99	0.55	1575.07	1252706	124	5
			Tool-calling	GPT-4o	0.00	0.00	0.00	1237.00	2489890	157	6
			ReAct	Claude 3.5	0.00	0.00	0.00	3123.35	1933814	47	0
			ReAct	GPT-4o	0.80	0.99	0.33	1188.14	1236142	104	1
		workflow	Tool-calling	Claude 3.5	0.67	0.89	0.44	1606.94	900033	120	3
			Tool-calling	GPT-4o	0.07	0.22	0.0001	1099.34	1549820	150	7
			ReAct	Claude 3.5	0.00	0.00	0.00	3081.22	1279999	47	0
		Subtask	ReAct	GPT-4o	0.53	0.97	0.06	1817.45	1329878	95	2
			Tool-calling	Claude 3.5	0.50	0.71	0.28	4098.50	4427260	362	79
			Tool-calling	GPT-4o	0.38	0.63	0.25	12957.03	15868903	629	40
			ReAct	Claude 3.5	0.28	0.36	0.25	3893.25	6921097	349	62
			ReAct	GPT-4o	0.52	0.82	0.27	3220.80	1996441	334	6
			Tool-calling	Claude 3.5	0.30	0.42	0.25	4638.77	6404970	342	81
			Tool-calling	GPT-4o	0.48	0.69	0.29	10539.12	9541567	566	27
			ReAct	Claude 3.5	0.30	0.42	0.25	6541.34	6996892	511	111
			ReAct	GPT-4o	0.53	0.84	0.29	3328.19	2892921	334	9
		workflow	Tool-calling	Claude 3.5	0.35	0.54	0.20	682.25	4714626	361	85
			Tool-calling	GPT-4o	0.48	0.74	0.28	13034.44	10529810	580	12
			ReAct	Claude 3.5	0.25	0.25	0.25	670.01	8800	97	7
			ReAct	GPT-4o	0.45	0.77	0.26	4450.11	5727793	455	102

References

- [1] Carlos E. Jimenez et al. “SWE-bench: Can Language Models Resolve Real-World GitHub Issues?” In: *arXiv preprint arXiv: 2310.06770* (2023).
- [2] Jun Shern Chan et al. “MLE-bench: Evaluating Machine Learning Agents on Machine Learning Engineering”. In: *International Conference on Learning Representations* (2025).
- [3] Alexander Novikov et al. “AlphaEvolve: A coding agent for scientific and algorithmic discovery”. In: *arXiv preprint arXiv: 2506.13131* (2025).
- [4] Juraj Gottweis et al. “Towards an AI co-scientist”. In: *arXiv preprint arXiv: 2502.18864* (2025).
- [5] Xingyao Wang et al. “Executable Code Actions Elicit Better LLM Agents”. In: *Forty-first International Conference on Machine Learning*, 2024. URL: <https://openreview.net/forum?id=jJ9BoXAffa>.
- [6] Zhengyao Jiang et al. “AIDE: AI-Driven Exploration in the Space of Code”. In: *arXiv preprint arXiv: 2502.13138* (2025).
- [7] Michael D. Skarlinski et al. “Language agents achieve superhuman synthesis of scientific knowledge”. In: *arXiv preprint arXiv: 2409.13740* (2024).
- [8] Yuan Chiang et al. “LLaMP: Large Language Model Made Powerful for High-fidelity Materials Knowledge Retrieval and Distillation”. In: *arXiv preprint arXiv: 2401.17244* (2024).
- [9] Shuxiang Cao et al. “Agents for self-driving laboratories applied to quantum computing”. In: *arXiv preprint arXiv: 2412.07978* (2024).
- [10] Quintina Campbell et al. “MDCrow: Automating Molecular Dynamics Workflows with Large Language Models”. In: *arXiv preprint arXiv: 2502.09565* (2025).
- [11] Yunheng Zou et al. “El Agente: An Autonomous Agent for Quantum Chemistry”. In: *arXiv preprint arXiv: 2505.02484* (2025).
- [12] Mhd Hussein Murtada, Z. Faidon Brotzakis, and Michele Vendruscolo. “MD-LLM-1: A Large Language Model for Molecular Dynamics”. In: *arXiv preprint arXiv: 2508.03709* (2025).
- [13] Indrajeet Mandal et al. “Autonomous Microscopy Experiments through Large Language Model Agents”. In: *arXiv preprint arXiv: 2501.10385* (2024).
- [14] Andres M. Bran et al. “Augmenting large language models with chemistry tools”. In: *Nature Machine Intelligence* 6.5 (May 2024), pp. 525–535. ISSN: 2522-5839. DOI: 10.1038/s42256-024-00832-8. URL: <http://dx.doi.org/10.1038/s42256-024-00832-8>.
- [15] Kourosh Darvish et al. “ORGANA: A robotic assistant for automated chemistry experimentation and characterization”. In: *Matter* 8.2 (Feb. 2025), p. 101897. ISSN: 2590-2385. DOI: 10.1016/j.matt.2024.10.015. URL: <http://dx.doi.org/10.1016/j.matt.2024.10.015>.
- [16] Daniil A. Boiko et al. “Autonomous chemical research with large language models”. In: *Nature* 624.7992 (Dec. 2023), pp. 570–578. ISSN: 1476-4687. DOI: 10.1038/s41586-023-06792-0. URL: <http://dx.doi.org/10.1038/s41586-023-06792-0>.
- [17] Hjalmar Wijk et al. “RE-Bench: Evaluating frontier AI R&D capabilities of language model agents against human experts”. In: *arXiv preprint arXiv: 2411.15114* (2024).
- [18] Megan Kinniment et al. “Evaluating Language-Model Agents on Realistic Autonomous Tasks”. In: *arXiv preprint arXiv: 2312.11671* (2023).
- [19] Zachary S. Siegel et al. “CORE-Bench: Fostering the Credibility of Published Research Through a Computational Reproducibility Agent Benchmark”. In: *arXiv preprint arXiv: 2409.11363* (2024).
- [20] Grégoire Mialon et al. “GAIA: a benchmark for General AI Assistants”. In: *arXiv preprint arXiv: 2311.12983* (2023).
- [21] Ludovico Mitchener et al. “BixBench: a Comprehensive Benchmark for LLM-based Agents in Computational Biology”. In: *arXiv preprint arXiv: 2503.00096* (2025).
- [22] Ziru Chen et al. “ScienceAgentBench: Toward Rigorous Assessment of Language Agents for Data-Driven Scientific Discovery”. In: *arXiv preprint arXiv: 2410.05080* (2024).
- [23] Jingyi Chai et al. “SciMaster: Towards General-Purpose Scientific AI Agents, Part I. X-Master as Foundation: Can We Lead on Humanity’s Last Exam?” In: *arXiv preprint arXiv: 2507.05241* (2025).

- [24] Minyang Tian et al. “SciCode: A Research Coding Benchmark Curated by Scientists”. In: *arXiv preprint arXiv: 2407.13168* (2024).
- [25] Hao Cui et al. “CURIE: Evaluating LLMs On Multitask Scientific Long Context Understanding and Reasoning”. In: *arXiv preprint arXiv: 2503.13517* (2025).
- [26] Siddharth Narayanan et al. “Aviary: training language agents on challenging scientific tasks”. In: *arXiv preprint arXiv: 2412.21154* (2024).
- [27] Microsoft. *AutoGen: A Programming Framework for Agentic AI*. <https://github.com/microsoft/autogen>. Accessed: 2025-08-20. 2025. URL: <https://github.com/microsoft/autogen>.
- [28] Sayash Kapoor et al. “AI Agents That Matter”. In: *Trans. Mach. Learn. Res.* 2025 (2025). URL: <https://openreview.net/forum?id=Zy4uFzMviZ>.
- [29] Nawaf Alampara, Mara Schilling-Wilhelmi, and Kevin Maik Jablonka. “Lessons from the trenches on evaluating machine-learning systems in materials science”. In: *arXiv preprint arXiv: 2503.10837* (2025).
- [30] Yue Huang et al. “MetaTool Benchmark for Large Language Models: Deciding Whether to Use Tools and Which to Use”. In: *arXiv preprint arXiv: 2310.03128* (2023).
- [31] Shunyu Yao et al. *τ -bench: A Benchmark for Tool-Agent-User Interaction in Real-World Domains*. 2024. arXiv: 2406.12045 [cs.AI]. URL: <https://arxiv.org/abs/2406.12045>.
- [32] Daan Frenkel and Berend Smit. *Understanding molecular simulation: from algorithms to applications*. Elsevier, 2023.
- [33] A. P. Thompson et al. “LAMMPS - a flexible simulation tool for particle-based materials modeling at the atomic, meso, and continuum scales”. In: *Comp. Phys. Comm.* 271 (2022), p. 108171. DOI: 10.1016/j.cpc.2021.108171.
- [34] Matthew K. Horton et al. “Accelerated data-driven materials science with the Materials Project”. In: *Nature Materials* (July 2025). ISSN: 1476-4660. DOI: 10.1038/s41563-025-02272-0. URL: <http://dx.doi.org/10.1038/s41563-025-02272-0>.
- [35] Lowik Chanussot et al. “Open catalyst 2020 (OC20) dataset and community challenges”. In: *Acs Catalysis* 11.10 (2021), pp. 6059–6072.
- [36] Tianqi Chen and Carlos Guestrin. “Xgboost: A scalable tree boosting system”. In: *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*. 2016, pp. 785–794.
- [37] Kevin Maik Jablonka, Luc Patiny, and Berend Smit. “Making Molecules Vibrate: Interactive Web Environment for the Teaching of Infrared Spectroscopy”. In: *Journal of Chemical Education* 99.2 (Jan. 2022), pp. 561–569. ISSN: 1938-1328. DOI: 10.1021/acs.jchemed.1c01101. URL: <http://dx.doi.org/10.1021/acs.jchemed.1c01101>.
- [38] Damiano Banfi and Luc Patiny. “www.nmrdb.org: Resurrecting and Processing NMR Spectra On-line”. In: *CHIMIA* 62.4 (Apr. 2008), p. 280. ISSN: 0009-4293. DOI: 10.2533/chimia.2008.280. URL: <http://dx.doi.org/10.2533/chimia.2008.280>.
- [39] Andrés M. Castillo, Luc Patiny, and Julien Wist. “Fast and accurate algorithm for the simulation of NMR spectra of large spin systems”. In: *Journal of Magnetic Resonance* 209.2 (Apr. 2011), pp. 123–130. ISSN: 1090-7807. DOI: 10.1016/j.jmr.2010.12.008. URL: <http://dx.doi.org/10.1016/j.jmr.2010.12.008>.
- [40] João Aires-de-Sousa, Markus C. Hemmer, and Johann Gasteiger. “Prediction of ¹H NMR Chemical Shifts Using Neural Networks”. In: *Analytical Chemistry* 74.1 (Dec. 2001), pp. 80–90. ISSN: 1520-6882. DOI: 10.1021/ac010737m. URL: <http://dx.doi.org/10.1021/ac010737m>.
- [41] Inc. Modal Labs. *Modal Python Package*. Version 1.0. Accessed: 2025-08-21. 2025. URL: <https://modal.com/docs/guide>.
- [42] Shyue Ping Ong et al. “Python Materials Genomics (pymatgen): A robust, open-source python library for materials analysis”. In: *Computational Materials Science* 68 (Feb. 2013), pp. 314–319. ISSN: 0927-0256. DOI: 10.1016/j.commatsci.2012.10.028. URL: <http://dx.doi.org/10.1016/j.commatsci.2012.10.028>.