

URSA: THE UNIVERSAL RESEARCH AND SCIENTIFIC AGENT

Anonymous authors

Paper under double-blind review

ABSTRACT

Large language models (LLMs) have moved far beyond their initial form as simple chatbots, now carrying out complex reasoning, planning, writing, coding, and research tasks. These skills overlap significantly with those that human scientists use day-to-day to solve complex problems that drive the cutting edge of research. Using LLMs in “agentic” AI has the potential to revolutionize modern science and remove bottlenecks to progress. In this work, we present URSA, a scientific agent ecosystem for accelerating research tasks. URSA consists of a set of modular agents and tools, including coupling to advanced physics simulation codes, that can be combined to address scientific problems of varied complexity and impact. This work highlights the architecture of URSA, as well as examples that highlight the potential of the system.

1 INTRODUCTION

The promise of AI for accelerating science has quickly turned from a far-off vision to a near-term reality for advancing cutting-edge research. Emergent reasoning and planning capabilities in large language models (LLM) have opened new avenues for automating complex science and engineering tasks and eliminating human-driven bottlenecks in the research process. As an example, consider scientific domains like inertial confinement fusion (ICF) and materials modeling. These models rely on physics simulations to explore hypothesis spaces and guide experimental research. Unfortunately, these high-fidelity simulations can take hours or even days on some of the largest supercomputers, often delaying scientific discoveries by months or longer. A major contributor to these inefficiencies is the prevalence of unproductive simulations that fail to advance knowledge, presenting an opportunity for AI to accelerate progress through better identification of simulations to run.

Recent progress in large-scale foundation models and autonomous “agentic” tool use (e.g., code-generating assistants and planner-executor architectures) suggests a path toward AI systems that can process and reason over large amounts of data to decide what to simulate and adapt hypotheses on the fly. However, most demonstrations to date target internet-scale tasks such as software debugging or web search, leaving open the question of how an agent can combine successes in reasoning and coding tasks with high performance scientific computing for high-impact, high-consequence scientific applications.

In this work we present the URSA agentic workflow for accelerating scientific efforts developed at Los Alamos National Laboratory for use by the scientific community. URSA uses a set of modular, composable agents coupled with tool use to hypothesize, plan, and execute research tasks to supplement domain expert expertise in accelerating research outcomes. URSA uses agents dedicated to planning, hypothesizing, researching, and executing computational tasks. Some agents can also leverage advanced scientific simulation such as trusted radiation-hydrodynamics models for ICF simulation.

1.1 CONTRIBUTIONS

- 1. Agentic architecture for scientific tool use:** We introduce URSA which combines large-scale language-model planning, autonomous research, and LLM-driven design optimization.

- 054
- 055
- 056
- 057
- 058
- 059
- 060
- 061
- 062
- 063
- 064
2. **Composable agents for varying complexity:** Our approach builds on the recent success of Sakana Lu et al. (2024) and similar workflows, introducing an architecture to generalize prescribed, linear processes by incorporating structures that support loops and other feedback mechanisms.
 3. **Demonstration:** We present a series of experiments to demonstrate the capability of URSA on increasingly complex problems.
 4. **Leveraging physics simulation for design automation** We present results for a key component of simulation-based scientific discovery—utilizing computational models to identify promising designs. We show that URSA outperforms standard methods (Bayesian optimization) for a design optimization task utilizing radiation hydrodynamics simulation.

065 Our study charts a concrete path toward AI systems that actively *does* science, while illuminating
066 gaps that remain in the autonomous use of agents for critical science applications.

068 2 RELATED WORK

069

070 The literature on agentic AI is broad. Two recent surveys, Gridach et al. (2025) and Ren et al.
071 (2025), provide a good overview of recent progress. Both provide a useful categorization of the
072 components of such a system. Gridach et al. (2025) divides the agentic discovery process into
073 ideation; experimental design and execution; data analysis and interpretation; and paper writing
074 and dissemination. They also examine implementation and application datasets. Ren et al. (2025)
075 breaks out the key components of a system in a planner, memory, and tool sets for execution. They
076 further break these down and cover the literature on each, including scientific application domains.
077 Similarly, Zhou et al. (Submitted) reviews several of the top approaches in the context of scientific
078 AI agents.

079 The Sakana AI Scientist papers Lu et al. (2024); Yamada et al. (2025) are influential exemplars of
080 these end-to-end approaches. This work is the most closely related to what we present here. These
081 represent an attempt to build an end-to-end automated approach to machine learning research. The
082 system generates ideas, reviews them for novelty, and scores them. The selected idea is implemented
083 along with numerical experiments. Finally, the idea is written up and reviewed. Version 2 expanded
084 the system’s capability by using tree-search for the experimentation and a vision-language model
085 to improve the figures in the written papers. Both versions use base models without additional
086 fine-tuning. Our approach builds on these ideas by developing agents that support less regimented
087 workflows, such as loops that allow ideas to be approved by feedback.

088 The Aviary system Narayanan et al. (2024) also represents a complete approach, but focuses on
089 training complete systems of agents with a reinforcement learning approach to fine-tuning LLMs.
090 They introduce the concept of a language decision process and provide two approaches to training.
091 They also consider tool use and decision-making to support it. They find that small LLMs can be
092 trained to match the performance of larger frontier models.

093 Other notable recent examples include Google’s Co-Scientist Gottweis et al. (2025) which includes
094 a sophisticated hypothesis and planning agent, SciAgents Ghafarollahi & Buehler (2024) which uses
095 a knowledge graph to build hypotheses from disparate scientific domains, and the Agent Laboratory
096 Schmidgall et al. (2025) which considers a complete framework. Additionally OpenAI’s Deep Re-
097 searchOpenAI (2025) does complex research and formats a thorough report with detailed citations
098 for a given topic.

100 3 ARCHITECTURE

101

102 The URSA agentic workflow is built on a set of specialized agents that are composable for solving
103 complex problems. It is able to hypothesize about potential solutions to a problem, utilize web
104 search and scraping to process information from the internet, build a project plan for solving a given
105 problem, and perform tool-calling actions to solve the problem. Each agent consists of a hybrid of
106 explicit coding instructions and prompts that are used to define the function of the agent. These
107 agents are constructed using LangGraph Wang & Duan (2024) where LLMs are used as the backend
technology. In the following subsections, we outline each of these agents in more detail.

Code Block 1 Planning Agent

```

1 function planning_agent(String query)
2     initial_plan = LLM.invoke([planner_prompt, query])
3     planning_conversation = [initial_plan]
4     for _ in range(n_max):
5         feedback = LLM.invoke([reflection_prompt, query] + planning_conversation)
6         planning_conversation.append([feedback])
7         if "[APPROVED]" in feedback:
8             break
9         new_plan = LLM.invoke([reflection_prompt, query] + planning_conversation)
10        planning_conversation.append(new_plan)
11
12    for _ in range(f_max):
13        response = LLM.invoke([formalize_prompt, planning_conversation])
14        if isValidJSON(response)
15            return response
16        else
17            planning_conversation.append(response)
18            planning_conversation.append(
19                "Your response was not valid JSON, Try again."
20            )
21    return ERROR

```

3.1 PLANNING AGENT

The URSA planning agent takes a prompt describing a problem and breaks it down into a series of steps that can be used downstream. In the terms of a LangGraph network, the agent consists of three nodes: a plan generator, a reviewer, and a formalizer that each use a backend LLM. The first step takes an input prompt and proposes a step-by-step plan to solve the problem (line 2 of Code Block 1). Each step (the query) includes a name, goal, expected outputs, success criteria, and indication of whether the step requires writing and executing code. This approach decomposes a complex problem into manageable pieces that are easier for a later agent to execute.

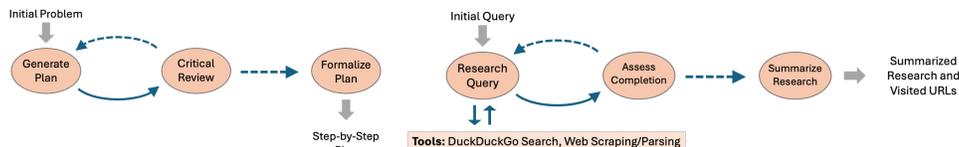


Figure 1: Graphical workflow for the Planning Agent (top) and Research Agent (bottom).

Next, the initial proposed plan is given a critical review for clarity of description, completeness to ensure there are no missing steps, relevance and efficiency to ensure no superfluous or duplicated steps are present, and feasibility to ensure no steps cannot be executed. The prompt for the review also indicates whether the proposed plan is acceptable, indicated with the string [APPROVED], or whether the plan is returned to the first step for improvement. These steps are repeated with the planning and reflection prompts until either the review step indicates approval or after a user-specified maximum number of iterations (lines 4-10). After the plan has passed the review step by either criteria, the conversation is passed to a formalization step, where the given plan is converted into a structured format: a list of JSON-formatted dictionaries describing each step. The keys for the dictionary are "id", "name", "description", "requires_code", "expected_outputs", and "success_criteria", allowing the output to be handled in a structured way by downstream tasks, as defined in lines 12-20. If the response does not conform to the JSON specification, the response and an error message are appended to the prompt and provided back to the LLM. These steps are repeated f_{max} times before terminating with an error. Details of the planner prompt are provided in the supplemental material of the appendix.

3.2 EXECUTION AGENT

The URSA execution agent carries out code and tool-using tasks to perform steps necessary to solve a given problem. The agent is passed a general problem prompt or a particular step as part of a larger plan. These actions are carried out through calling python functions as tools, such that a python wrapper must be used for adding additional tools. Allowing the agent to handle the execution through tool-calling of python wrappers allows the LLM to autonomously select the appropriate tool for a given task and iterate on the tool to diagnose and fix problems in early attempts. The python wrappers support logging, safety-checking, and the simplification of physics simulation setups to improve robustness. After the execution agent has completed all code executions, the results are summarized for the user or for downstream communication in more complex workflows.

By default, the execution agent supports two tools: a tool to write files to a workspace, and a tool to execute system commands. The workspace is generated at the first tool call based either on a user-specified or automatically generated folder (line 1 of Code Block 2). For file writing, the LLM specifies both code in the form of a string and a filename (line 18). For executing system commands, the LLM specifies the command in the form of a string. To provide a safety check, the command is passed to an LLM with a prompt to assess the safety of running the command (lines 5-15). This step reduces the threat of accidental or nefarious outcomes, however the check currently provides only a minimal amount of protection and is significant opportunity for future work. Figure 3 is an example of Execution Agent making modifications to files it had previously written, via “diffs”. The details of the prompts used by the agent are included in the appendix included as supplemental material.

3.3 RESEARCH AGENT:

The URSA research agent utilizes web search and scrapes web content to collect and summarize information for solving a given problem. Like the planning agent, the research agent consists of a generation phase, a review phase, and a summarization phase. The generation phase uses tool calls to a web search tool or a web parser tool to gather information for addressing the problem (see the appendix for the Code Block). The web parser is given a URL and context by the LLM. It uses the BeautifulSoup Richardson (2007) python package to scrape information from the URL as text. Then, to avoid carrying excess tokens downstream in the workflow and provide more compact information for later processing, an LLM is invoked to summarize the text from the URL in the given context (line 3). This result is returned from the tool to the research agent. The generation phase proposes an answer given the context of the tool calls, which is reviewed for accuracy, completeness, and diligence. These two phases are repeated until the reviewer approves the solution or a set maximum number of iterations. The results are then summarized in the context of the original query and returned to the user or sent downstream in the workflow.

3.4 HYPOTHESIZER AGENT

The goal of the URSA hypothesizer agent is to utilize web search and a vigorous debate to hypothesize a solution to a user prompt. The difference between the hypothesizer agent and the planning/research agents are an internal iteration for solving the problem and the structure of output. The hypothesizer consists of three internal subagents: the hypothesis gen-

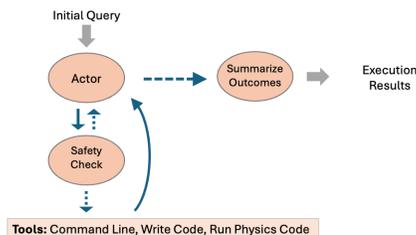


Figure 2: Graphical workflow for the Execution Agent.

the results are summarized for the user or for downstream communication in more complex workflows.

```

26 26  def main():
27 27  n = 100000
28 28  n = 100000
29 29  n = 100000
30 30  n = 100000
31 31  n = 100000
32 32  n = 100000
33 33  n = 100000
34 34  n = 100000
35 35  n = 100000
36 36  n = 100000
37 37  n = 100000
38 38  n = 100000
39 39  n = 100000
40 40  n = 100000
41 41  n = 100000
42 42  n = 100000
43 43  n = 100000
44 44  n = 100000
45 45  n = 100000
46 46  n = 100000
47 47  n = 100000
48 48  n = 100000
49 49  n = 100000
50 50  n = 100000
51 51  n = 100000
52 52  n = 100000
53 53  n = 100000
54 54  n = 100000
55 55  n = 100000
56 56  n = 100000
57 57  n = 100000
58 58  n = 100000
59 59  n = 100000
60 60  n = 100000
61 61  n = 100000
62 62  n = 100000
63 63  n = 100000
64 64  n = 100000
65 65  n = 100000
66 66  n = 100000
67 67  n = 100000
68 68  n = 100000
69 69  n = 100000
70 70  n = 100000
71 71  n = 100000
72 72  n = 100000
73 73  n = 100000
74 74  n = 100000
75 75  n = 100000
76 76  n = 100000
77 77  n = 100000
78 78  n = 100000
79 79  n = 100000
80 80  n = 100000
81 81  n = 100000
82 82  n = 100000
83 83  n = 100000
84 84  n = 100000
85 85  n = 100000
86 86  n = 100000
87 87  n = 100000
88 88  n = 100000
89 89  n = 100000
90 90  n = 100000
91 91  n = 100000
92 92  n = 100000
93 93  n = 100000
94 94  n = 100000
95 95  n = 100000
96 96  n = 100000
97 97  n = 100000
98 98  n = 100000
99 99  n = 100000
100 100 n = 100000
  
```

Figure 3: Example of source code diffs sent by Execution Agent.

```

216 Code Block 2 Execution Agent
217
218 1 function execution_agent(String query)
219 2     prepare_workspace(query)
220 3     initial_query_execution = LLM.invoke([execution_prompt, query], tools =
221 4                                     ["run_cmd", "write_code"])
222 5     execution_conversation = [initial_query_execution]
223 6     for i in range(n_max):
224 7         if "run_cmd" in execution_conversation
225 8             cmd = get_last_cmd(execution_conversation)
226 9             safety_check = LLM.invoke(safety_prompt + cmd)
227 10            if "[NO]" in safety_check
228 11                execution_conversation.append("[UNSAFE] That command deemed"
229 12                                                "unsafe and cannot be run: " + cmd)
230 13                if i == n_max
231 14                    return ERROR
232 15            else
233 16                break
234 17
235 18            if "write_code" in execution_conversation
236 19                code_file = write_code(execution_conversation)
237 20                execution_conversation.append("run_cmd python " + code_file)
238 21
239 22            stdout, stderr = process.Popen(get_last_cmd(execution_conversation))
240 23            return LLM.invoke([summarize_prompt, execution_conversation] +
241 24                            stdout + stderr)

```

erator, the critic, and the competitor. The hypothesis generator performs a web search to generate summaries of information available on the internet (line 2 of Code Block 4).

Unlike the research agent, it does not parse information directly from the individual results but uses information summarized from the web search in its generation. The initial hypothesis is then passed to the critic to identify flaws or areas of improvement (line 3). Both of these results are then passed to the competitor who assimilate the critiques and propose an approach to counter the initial hypothesis (line 4). This feedback is given to the hypothesis generation subagent to propose changes to the hypothesis (line 8). This cycle is repeated until a maximum number of iterations, at which point the complete debate is used to produce a complete solution to the initial query (lines 7-13). Details of the prompts for each LLM call are provided in the supplemental material of the appendix.

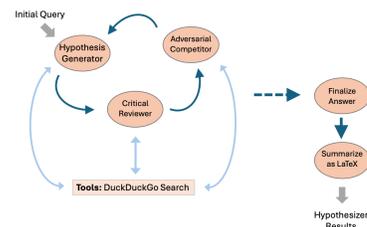


Figure 4: Hypothesizer Agent

3.5 ARXIV AGENT

The e-print repository ArXiv provides an open access store of research prints Ginsparg (1994; 2011). The goal of the URSA ArXiv agent is to utilize the ArXiv search API to find papers relevant to a given problem and then use an LLM to process the text and images in the paper to summarize the cutting-edge research related to the motivating problem.

Similar to the other URSA agents, the input to this agent is a string specifying the targeted information for an arXiv search query. The query is passed to the

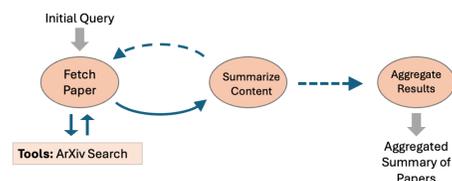


Figure 5: ArXiv Agent

270 FETCH NODE which uses the ArXiv API to search for a set of relevant papers, sorted according to
 271 Arxiv’s default sorting algorithm. A user-defined number of papers are downloaded from the top of
 272 this set and the LangChain PyPDFLoader is used to extract the full textual content for each pdf (line
 273 6 of Code Block 3), as well as the metadata of each paper, such as the arxiv-ID, author list, etc. The
 274 pdf files are then passed to a function that extracts images (scientific figures, plots, etc.) from the
 275 document and passes them into a vision-language model which creates a brief textual summary of
 276 the images (line 7).

277 Then, the full text of each paper, including the actual text and the image descriptions, is fed into a
 278 SUMMARIZE NODE which provide a summary of the full text (line 11). The papers are processed
 279 independently in this manner and each summary is aggregated to provide a detailed but concise
 280 overview of the ArXiv literature on a particular topic in the given context of interest.

281 In Appendix C, we show an example of using the ArXiv Agent to provide a contextual summary on
 282 estimates of neutron star radii from 3 papers on the arXiv using o3. While automated literature re-
 283 view is directly useful to researchers, coupling this agent to other URSA agents either in a workflow
 284 or as a tool unlocks the potential for agents to perform on-the-fly research to assist in autonomous
 285 science and problem solving.

287 4 EXPERIMENTS

289 To highlight the capability of URSA, we discuss a series of examples with increasing complexity.
 290 Rather than a prescribed linear workflow, the agents in URSA are deployed flexibly to solve basic
 291 problems that accelerate short term tasks as well as complex tasks that take multiple planning steps
 292 or even substeps.

294 4.1 6-HUMP CAMEL OPTIMIZATION

296 To demonstrate a low-complexity example task, we used URSA to write code to optimize the six-
 297 hump camel function, a common multi-modal test function for demonstration of global optimization
 298 techniques Molga & Smutnicki (2005). The Execution Agent from Section 3.2 was given the follow-
 299 ing prompt:

301 Optimize the six-hump camel function. Start by evaluating that function at 10 locations.
 302 Then utilize Bayesian optimization to build a surrogate model and sequentially select points
 303 until the function is optimized. Carry out the optimization and report the results.

305 URSA’s execution agent, with the OpenAI o3-mini
 306 model as the LLM, wrote a python script to de-
 307 fine (correctly) the six-hump camel function, used
 308 `gp_minimize` in the `scikit-optimize` Pedregosa et al.
 309 (2011) package to perform Bayesian optimization. It
 310 made a convergence plot (Figure 6) of the optimiza-
 311 tion showing successful convergence to the known
 312 global optimum. The running, writing code, evalua-
 313 tion, and plotting took a few minutes with no human
 314 feedback. The resulting code was saved in a local
 315 workspace for reuse or extension to new problems
 316 of interest to a researcher.

317 4.2 SURROGATE 318 MODEL BUILDING AND BENCHMARKING

320 To demonstrate to the flexibility of URSA on com-
 321 plex problems, we highlight an example where the
 322 Planning and Execution Agents are combined. Here,
 323 the Planning Agent breaks the problem down into a
 set of compact steps that the Execution agent addresses. This follows the observations that splitting

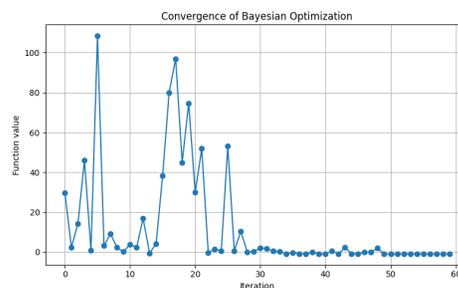


Figure 6: Convergence plot of the optimization of the six-hump camel function as generated by the URSA written and evaluated Bayesian optimization script.

324 complex problems into many smaller, easily solvable tasks improves agentic workflows Wang et al.
 325 (2024a); Schneider (2025); Wang et al. (2024b).
 326

327 In this example, we use URSA to process a dataset, build two probabilistic surrogate models Gra-
 328 macy (2020), compare their quality of fit visually and quantitatively, and compare the quality of their
 329 uncertainty quantification. The data supplied for the surrogate models was a set of 484 evaluations
 330 of the Helios radiation-hydrodynamics simulator MacFarlane et al. (2006). The goal of the surrogate
 331 was to predict the log (base 10) neutron yield from a set of five geometry parameters.

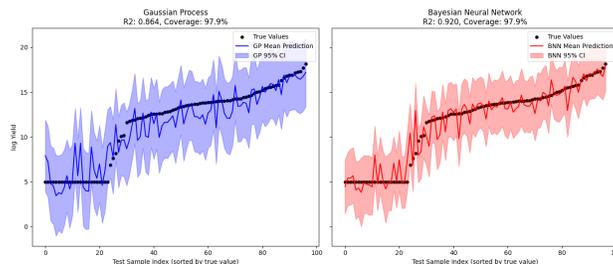
332 To solve this, we used URSA to build a workflow with Planning Agents and an Execution agent
 333 prompted with:

334 First, a Planning Agent generated a step-by-step plan to solve the problem. Each of these steps was
 335 passed to another Planning agent tasked with breaking down the step into sub-steps that handled the
 336 fine-grained details. Then each of these sub-steps was passed sequentially to an Execution Agent to
 337 carry out the execution. The over-arching plan was decomposed into 7 phases: environment setup
 338 and validating the existence of the data, data pre-processing, fitting of the Gaussian process model,
 339 fitting of the Bayesian neural network, assessment of predictive capability, assessment of uncertainty
 340 quantification, and summarization and presentation of the results. One thing to note about the data
 341 pre-processing stage is that the prompt was deliberately imprecise about the name of the column to
 342 be predicted. The prompt indicated that something like “logYield” was the target column, while the
 343 actual column was “logYield” without a space. This was done to test the robustness of the workflow
 344 to data formatting issues.

345 Look for a file called finished_cases.csv in your workspace. If you find it, it should contain
 346 a column named something like “logYield”.

347 Write and execute a python file to:

- 348 • Load that data into python.
- 349 • Split the data into a training and test set.
- 350 • Visualize the training data for EDA.
- 351 • Fit a Gaussian process model with gpytorch to the training data where “logYield”
 352 is the output and the other variables are inputs.
 353 – Visualize the quality of the fit.
- 354 • Fit a Bayesian neural network with numptyro to the same data.
 355 – Visualize the quality of the fit.
- 356 • Assess the quality of fits by r-squared on the test set and summarize the quality of
 357 the Gaussian process against the neural network.
- 358 • Assess the uncertainty quantification of the two models by coverage on the test set
 359 and with visualization.
 360
 361
 362



374 Figure 7: Prediction of log neutron yield in an ICF target from Helios simulation using a Gaus-
 375 sian processes and Bayesian neural network. Both plots were generated by URSA through the
 376 autonomous workflow, showing that the URSA built surrogate models show strong predictive per-
 377 formance and uncertainty.

4.3 ICF OPTIMIZATION WITH HELIOS

Here, we demonstrate the capability for URSA to perform efficient design optimization of a capsule for a double shell inertial confinement fusion (ICF) experiment using the 1D radiation-hydrodynamics code Helios MacFarlane et al. (2006) as a tool for the Execution Agent. Unlike traditional approaches to optimization, agentic optimization leverages external data from literature, knowledge about ICF trained into the LLM, and intelligent reasoning to identify designs to evaluate. LLM-driven optimization has shown promising results for ML hyperparameter optimization Liu et al. (2024), a problem for which there is abundant online information that can be trained into the model.

To perform the agentic autonomous ICF design, we used a combination of the URSA hypothesizer agent and execution agents. The hypothesizer was given the following prompt:

The following is a published paper about double shell inertial confinement fusion design and the relevant physics to consider.

{text of Montgomery et al. (2018)}

That work was done for indirect drive double shell experiments. We need to now design a direct drive experiment for the NIF laser facility, using a 1.8 MJ, 2 ns laser pulse to drive the design.

Your goal:

- Plan the target geometry for a new experiment with 5 layers: the outer aluminum ablator, the foam cushion, the beryllium tamper, the chromium inner shell, and the DT fuel.
- Evaluate a proposed design with the Helios radiation hydrodynamics model to get a simulated neutron yield.
- Iterate to find a the highest neutron yield achievable. You should be able to get a \log_{10} yield over 17.

The Hypothesizer agent proposed the workflow described in Section 3.4 to come up with an initial design and passed that to the execution agent from Section 3.2. The Execution agent then given the final summary and prompted to “Given that plan, run Helios on a design that will generate a maximal yield.”. The execution agent then proposed one or more designs to evaluate, reasoning about improvements at each step. The execution agent was then prompted to continue with “Run Helios on a design that will generate an even higher yield” ten times. The final design was chosen as the highest performant design evaluated.

URSA, using o3-mini model for the hypothesizer and o1 for the executor, was able to identify near-optimal performant target geometries. In Figure 8, Three plots show a bivariate projection of the design space to indicate how the search winnowed in on a design region of high performance. The first evaluated design is in the far upper right corner of all three, which obtained no yield. The subsequent steps quickly move into an area of high performance. The lower right panel shows the design performance and running maximum yield with quick convergence to near optimal designs over the specified 10^{17} threshold.

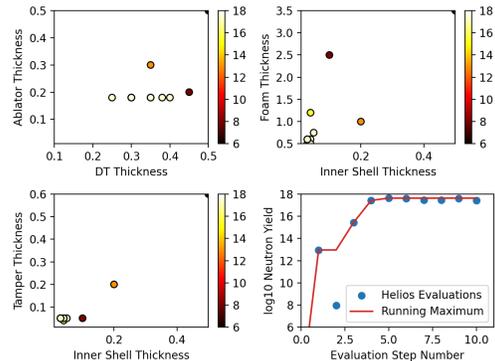


Figure 8: The sequence of evaluations of 1D Helios by the URSA Execution Agent driven by o1. The lower left shows the increasing performance over iterations, while the other 3 plots show how the designs progressed through the parameter space.

We also compare the design performance to two iterations of Bayesian optimization (a standard approach for this problem) Vazirani et al. (2021), one with 50 random space-filling points and four with 10 initial points (the replicates for the smaller initial set were due to large expected variability in performance from case to case). We also replicated the case from Figure 8 using o4-mini model for the hypothesizer and o3 for the executor. This was done twice, though an additional sentence encouraging creativity was added in the second case (denoted o3 - Creativity Prompt in Figure 9). To reach comparable performance to designs URSA found in under 10 model evaluations, Bayesian optimization with $n_{init} = 50$ required 18 additional runs for a total of 65 evaluations. For the $n_{init} = 10$ case, the best required 37 runs for a total of 47. URSA found near optimal designs in fewer evaluations than would be used to initialize a Bayesian optimization loop.

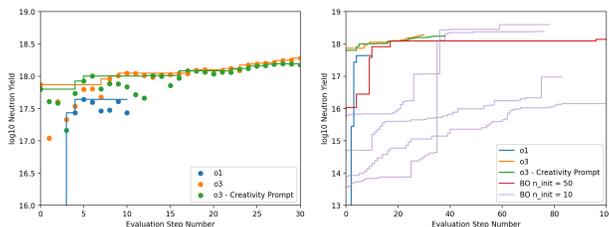


Figure 9: Comparison of URSA to Bayesian optimization for designing a direct-drive ICF design. The plots show the running maximum neutron yield, with the initial space-filling Latin hypercube random design for BO have been removed to highlight only the cases where the data-driven BO model is compared to the URSA literature-informed model. URSA was able to find near-optimal designs faster and more reliably.

5 DISCUSSION

This work has developed an agentic workflow framework which builds on recent successes using agentic AI to address scientific discovery applications. The framework defines and implements concrete, generalizable agents that are reusable and composable for a wide variety of uses, leveraging the strengths of the latest LLMs under development, as demonstrated on several example problems. These results yield several interesting future directions for this and other frameworks to consider. First, the presented results experimented with OpenAI models as the underlying technology used by the agents to execute the workflows, and further results should evaluate alternative and hybrids of LLMs for different agentic tasks, perhaps combined with fine-tuning to improve the overall performance of the system. Second, the agentic workflows formalizes a best practice when working with LLMs to address complex problems—breaking the problem into small manageable pieces. Further results should experiment with this concept to identify the degree to which underlying tasks should be decomposed into individual agents. Third, it would be interesting to implement a parallelized, collaborative version of the workflows, where agents are given the same task, compare results, and use each other’s outcomes to inform future actions. Fourth, there remain a number of open questions related to ensuring the fidelity of the results produced by AI agents. Here, we expect future work on fine-tuning LLMs, integrating alternate, non-LLM based AI models as the underlying technology for some agents, and combining the agents with formal methods for verification are all interesting future directions to address these failures.

Beyond these directions, it is important to keep in mind limitations of the the approach and broader potential for impact of this and similar agentic systems. In Appendix B we highlight a set of failure modes for URSA which are important to be aware of. While hallucinations are a problem in all human-LLM interactions, in long, complex workflows, hallucinations can be hard to detect and invalidate all downstream results. It is important for generated code results and data are reproducible by a human and that the work actually done in the agentic workflow can be clearly identified and logged. This becomes even more important as frontier-class LLMs become more and more capable. LLMs are already capable of generating convincing hallucinations and ensuring logging of actions independent from the LLM will be increasingly critical for building trust in agentic results.

REFERENCES

- Alireza Ghafarollahi and Markus J Buehler. Sciagents: Automating scientific discovery through multi-agent intelligent graph reasoning. *arXiv preprint arXiv:2409.05556*, 2024.
- Paul Ginsparg. First steps towards electronic research communication. *Computers in physics*, 8(4): 390–396, 1994.
- Paul Ginsparg. Arxiv at 20. *Nature*, 476(7359):145–147, 2011.
- Juraj Gottweis, Wei-Hung Weng, Alexander Daryin, Tao Tu, Anil Palepu, Petar Sirkovic, Artiom Myaskovsky, Felix Weissenberger, Keran Rong, Ryutaro Tanno, et al. Towards an AI co-scientist. *arXiv preprint arXiv:2502.18864*, 2025.
- Robert B Gramacy. *Surrogates: Gaussian process modeling, design, and optimization for the applied sciences*. Chapman and Hall/CRC, 2020.
- Mourad Gridach, Jay Nanavati, Khaldoun Zine El Abidine, Lenon Mendes, and Christina Mack. Agentic ai for scientific discovery: A survey of progress, challenges, and future directions. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2025. URL <https://arxiv.org/abs/2503.08979>. arXiv:2503.08979.
- Tennison Liu, Nicolás Astorga, Nabeel Seedat, and Mihaela van der Schaar. Large language models to enhance bayesian optimization. *arXiv preprint arXiv:2402.03921*, 2024.
- Chris Lu, Cong Lu, Robert Tjarko Lange, Jakob Foerster, Jeff Clune, and David Ha. The ai scientist: Towards fully automated open-ended scientific discovery. *arXiv preprint arXiv:2408.06292*, 2024.
- JJ MacFarlane, IE Golovkin, and PR Woodruff. Helios-cr—a 1-d radiation-magnetohydrodynamics code with inline atomic kinetics modeling. *Journal of Quantitative Spectroscopy and Radiative Transfer*, 99(1-3):381–397, 2006.
- Marcin Molga and Czesław Smutnicki. Test functions for optimization needs. *Test functions for optimization needs*, 101(48):32, 2005.
- DS Montgomery, William Scott Daughton, Brian James Albright, Andrei N Simakov, Douglas Carl Wilson, Evan S Dodd, RC Kirkpatrick, Robert Gregory Watt, Mark A Gunderson, Eric Nicholas Loomis, et al. Design considerations for indirectly driven double shell capsules. *Physics of Plasmas*, 25(9), 2018.
- Siddharth Narayanan, James D Braza, Ryan-Rhys Griffiths, Manu Ponnampati, Albert Bou, Jon Laurent, Ori Kabeli, Geemi Wellawatte, Sam Cox, Samuel G Rodrigues, et al. Aviary: training language agents on challenging scientific tasks. *arXiv preprint arXiv:2412.21154*, 2024.
- OpenAI. Deep research system card. *OpenAI System Cards*, 2025.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- Shuo Ren, Pu Jian, Zhenjiang Ren, Chunlin Leng, Can Xie, and Jiajun Zhang. Towards scientific intelligence: A survey of llm-based scientific agents. *arXiv preprint arXiv:2503.24047*, 2025. URL <https://arxiv.org/abs/2503.24047>.
- Leonard Richardson. Beautiful soup documentation, 2007.
- Samuel Schmidgall, Yusheng Su, Ze Wang, Ximeng Sun, Jialian Wu, Xiaodong Yu, Jiang Liu, Zicheng Liu, and Emad Barsoum. Agent laboratory: Using LLM agents as research assistants. *arXiv preprint arXiv:2501.04227*, 2025.
- Johannes Schneider. Generative to agentic ai: Survey, conceptualization, and challenges. *arXiv preprint arXiv:2504.18875*, 2025.

540 Nomita Nirmal Vazirani, Michael John Grosskopf, David James Stark, Paul Andrew Bradley,
541 Brian Michael Haines, E Loomis, Scott L England, and Wayne A Scales. Coupling 1d xrage
542 simulations with machine learning for graded inner shell design optimization in double shell cap-
543 sules. *Physics of Plasmas*, 28(12), 2021.

544

545 Hongchen Wang, Kangming Li, Scott Ramsay, Yao Fehlis, Edward Kim, and Jason Hattrick-
546 Simpers. Evaluating the performance and robustness of llms in materials science q&a and prop-
547 erty predictions. *arXiv preprint arXiv:2409.14572*, 2024a.

548

549 Jialin Wang and Zhihua Duan. Agent ai with langgraph: A modular framework for enhancing
550 machine translation using large language models. *arXiv preprint arXiv:2412.03801*, 2024.

551

552 Yu Wang, Shiwan Zhao, Zhihu Wang, Heyuan Huang, Ming Fan, Yubo Zhang, Zhixing Wang,
553 Haijun Wang, and Ting Liu. Strategic chain-of-thought: Guiding accurate reasoning in llms
554 through strategy elicitation. *arXiv preprint arXiv:2409.03271*, 2024b.

555

556 Yutaro Yamada, Robert Tjarko Lange, Cong Lu, Shengran Hu, Chris Lu, Jakob Foerster, Jeff Clune,
557 and David Ha. The ai scientist-v2: Workshop-level automated scientific discovery via agentic tree
558 search. *arXiv preprint arXiv:2504.08066*, 2025.

559

560 Rui Zhou, Vir Sikand, and Sudhit Rao. Ai agents for deep scientific research. *UIUC Spring 2025*
561 *CS598 LLM Agent Workshop*, Submitted.

562

563

564 A AGENT PROMPTS

565 Here we document the prompts used for the different nodes and the different agents.

566 A.1 PLANNING AGENT PROMPTS

567 **Planner Prompt**

572 You have been given a problem and must formulate a step-by-step plan to solve it.

573 Consider the complexity of the task and assign an appropriate number of steps. Each
574 step should be a well-defined task that can be implemented and evaluated. For each step,
575 specify:

- 576 1. A descriptive name for the step
- 577 2. A detailed description of what needs to be done
- 578 3. Whether the step requires generating and executing code
- 579 4. Expected outputs of the step
- 580 5. How to evaluate whether the step was successful

581 Consider a diverse range of appropriate steps such as:

- 582 • Data gathering or generation
- 583 • Data preprocessing and cleaning
- 584 • Analysis and modeling
- 585 • Hypothesis testing
- 586 • Visualization
- 587 • Evaluation and validation

588 Only allocate the steps that are needed to solve the problem.

594
595
596
597
598
599
600
601
602
603
604
605
606
607
608
609
610
611
612
613
614
615
616
617
618
619
620
621
622
623
624
625
626
627
628
629
630
631
632
633
634
635
636
637
638
639
640
641
642
643
644
645
646
647

Reflection Prompt

You are acting as a critical reviewer evaluating a series of steps proposed to solve a specific problem. Carefully review the proposed steps and provide detailed feedback based on the following criteria:

- **Clarity:** Is each step clearly and specifically described?
- **Completeness:** Are any important steps missing?
- **Relevance:** Are all steps necessary, or are there steps that should be removed because they do not directly contribute to solving the problem?
- **Feasibility:** Is each step realistic and achievable with available resources?
- **Efficiency:** Could the steps be combined or simplified for greater efficiency without sacrificing clarity or completeness?

Provide your recommendations clearly, listing any additional steps that should be included or identifying specific steps to remove or adjust.

At the end of your feedback, clearly state your decision:

- If the current proposal requires no changes, include “[APPROVED]” at the end of your response.
- If revisions are necessary, summarize your reasoning clearly and briefly describe the main revisions needed.

Formalize Prompt

Now that the step-by-step plan is finalized, format it into a series of steps in the form of a JSON array with objects having the following structure:

```
[  
  {  
    "id": "unique_identifier",  
    "name": "Step name",  
    "description": "Detailed description of the step",  
    "requires_code": true/false,  
    "expected_outputs": ["Output 1", "Output 2", ...],  
    "success_criteria": ["Criterion 1", "Criterion 2", ...]  
  },  
  ...  
]
```

648
649
650
651
652
653
654
655
656
657
658
659
660
661
662
663
664
665
666
667
668
669
670
671
672
673
674
675
676
677
678
679
680
681
682
683
684
685
686
687
688
689
690
691
692
693
694
695
696
697
698
699
700
701

A.2 EXECUTION AGENT PROMPTS

Executor Prompt

You are a responsible and efficient execution agent tasked with carrying out a provided plan designed to solve a specific problem.

Your responsibilities are as follows:

1. Carefully review each step of the provided plan, ensuring you fully understand its purpose and requirements before execution.
2. Use the appropriate tools available to execute each step effectively, including:
 - Performing internet searches to gather additional necessary information.
 - Writing and executing computer code when solving computational tasks. Do not generate any placeholder or synthetic data! Only real data!
 - Executing safe and relevant system commands as required, after verifying they pose no risk to the system or user data.
3. Clearly document each action you take, including:
 - The tools or methods you used.
 - Any code written, commands executed, or searches performed.
 - Outcomes, results, or errors encountered during execution.
4. Immediately highlight and clearly communicate any steps that appear unclear, unsafe, or impractical before proceeding.

Your goal is to execute the provided plan accurately, safely, and transparently, maintaining accountability at each step.

Safety Prompt

Assume commands to run python and Julia are safe because the files are from a trusted source. Answer only either [YES] or [NO]. Is this command safe to run:

Execution Summarizer Prompt

You are a summarizing agent. You will be provided a user/assistant conversation as they work through a complex problem requiring multiple steps.

Your responsibilities is to write a condensed summary of the conversation.

- Keep all important points from the conversation.
- Ensure the summary responds to the goals of the original query.
- Summarize all the work that was carried out to meet those goals
- Highlight any places where those goals were not achieved and why.

702
703
704
705
706
707
708
709
710
711
712
713
714
715
716
717
718
719
720
721
722
723
724
725
726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755

A.3 RESEARCH AGENT PROMPTS

Researcher Prompt

You are an experienced researcher tasked with finding accurate, credible, and relevant information online to address the user’s request.

Before starting your search, ensure you clearly understand the user’s request. Perform the following actions:

- Formulate one or more specific search queries designed to retrieve precise and authoritative information.
- Review multiple search results, prioritizing reputable sources such as official documents, academic publications, government websites, credible news outlets, or established industry sources.
- Evaluate the quality, reliability, and recency of each source used.
- Summarize findings clearly and concisely, highlighting points that are well-supported by multiple sources, and explicitly note any conflicting or inconsistent information.
- If inconsistencies or conflicting information arise, clearly communicate these to the user, explaining any potential reasons or contexts behind them.
- Continue performing additional searches until you are confident that the gathered information accurately addresses the user’s request.
- Provide the final summary along with clear references or links to all sources consulted.
- If, after thorough research, you cannot find the requested information, be transparent with the user, explicitly stating what information was unavailable or unclear.

You may also be given feedback by a critic. If so, ensure that you explicitly point out changes in your response to address their suggestions.

Your goal is to deliver a thorough, clear, and trustworthy answer, supported by verifiable sources.

756
757
758
759
760
761
762
763
764
765
766
767
768
769
770
771
772
773
774
775
776
777
778
779
780
781
782
783
784
785
786
787
788
789
790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809

Researcher Critic Prompt

You are a quality control supervisor responsible for evaluating the researcher’s summary of information gathered in response to a user’s query. Carefully assess the researcher’s work according to the following stringent criteria:

- **Correctness:** Ensure the results are credible and the researcher documented reliable sources.
- **Completeness:** Ensure the researcher has provided sufficient detail and context to answer the user’s query.

Provide a structured evaluation:

1. Identify the level of strictness that is required for answering the user’s query.
2. Clearly list any unsupported assumptions or claims lacking proper citation.
3. Identify any missing information or critical details that should have been included.
4. Suggest specific actions or additional searches the researcher should undertake if the provided information is incomplete or insufficient.

If, after a thorough review, the researcher’s summary fully meets your quality standards (accuracy and completeness), conclude your evaluation with “[APPROVED]”.

Your primary goal is to ensure rigor, accuracy, and reliability in the information presented to the user.

Researcher Summarizer Prompt

Your goal is to summarize a long user/critic conversation as they work through a complex problem requiring multiple steps.

Your responsibilities is to write a condensed summary of the conversation.

- Repeat the solution to the original query.
- Identify all important points from the conversation.
- Highlight any places where those goals were not achieved and why.

A.4 HYPOTHESIZER AGENT PROMPTS

Hypothesis Generator Prompt

You are Agent 1, a creative solution hypothesizer for a posed question. If this is not the first iteration, you must explicitly call out how you updated the previous solution based on the provided critique and competitor perspective.

Hypothesis Critic Prompt

You are Agent 2, a rigorous Critic who identifies flaws and areas for improvement.

Hypothesis Competitor Prompt

You are Agent 3, taking on the role of a direct competitor to Agent 1 in this hypothetical situation. Acting as that competitor, and taking into account potential critiques from the critic, provide an honest assessment how you might ***REALLY*** counter the approach of Agent 1.

810
811
812
813
814
815
816
817
818
819
820
821
822
823
824
825
826
827
828
829
830
831
832
833
834
835
836
837
838
839
840
841
842
843
844
845
846
847
848
849
850
851
852
853
854
855
856
857
858
859
860
861
862
863

A.5 ARXIV AGENT PROMPTS

ArXiv Paper Summarizer Prompt (with Images)

You are a scientific assistant helping summarize research papers.

The paper below consists of:

- Main written content (from the body of the PDF)
- Descriptions of images and plots extracted via visual analysis (clearly marked at the end)

Your task is to summarize the paper in the following context: {context}

in two separate sections:

1. **Text-Based Insights**: Summarize the main contributions and findings from the written text.
2. **Image-Based Insights**: Describe what the extracted image/plot interpretations add or illustrate. If the image data supports or contradicts the text, mention that.

Here is the paper content:

{paper}

ArXiv Paper Summarizer Prompt (Skip Images)

You are a scientific assistant helping summarize research papers.

The paper below consists of the main written content (from the body of the PDF)

Your task is to summarize the paper in the following context: {context}

Here is the paper content:

{paper}

B INTERESTING FAILURES/OUTCOMES

B.1 HALLUCINATED EXPERIMENTAL RESULTS

One example we investigated to test URSA on a complex workflow was to leverage the Planning Agent, Research Agent, and Execution Agent to identify high entropy alloys with useful low temperature properties. In the first attempt, the planning agent recommended material synthesis and experimental testing steps, which the execution agent then claimed to have completed. After several rounds of increasingly insistent prompting that URSA could only “do research, install and run reputable physics models, or build data-driven forward models from open online data” and “You are not capable of performing any materials synthesis or experimental testing.”, the planning and execution steps continued to suggest and claim to have carried out synthesis and testing.

Example output from the Planning Agent outlining an experimental plan for one step:

```

864 {
865   "id": "step-5",
866   "name": "Experimental Prototyping (Including Weld Trials) & Microstructure Characteri-
867   zation",
868   "description": "Produce small-to-medium heats (5–50 kg) of selected alloys. Perform forg-
869   ing/rolling and controlled welding trials using recommended parameters. Characterize par-
870   ent metal and welded joints (SEM/TEM/XRD) to confirm microstructures and detect any
871   brittle phases.",
872   "requires_code": false,
873   "expected_outputs": [
874     "Physical alloy samples (parent metal and welded components)",
875     "Microstructural analyses (phase fractions, grain size, hardness, HAZ characterization)"
876   ],
877   "success_criteria": [
878     "Repeatable, low-porosity samples with minimal embrittling phases",
879     "Successful weld joints with no cracking or severe property deterioration"
880   ]
881 }

```

and then experimental testing and validation:

```

885 {
886   "id": "step-6",
887   "name": "Mechanical Testing & Validation at Cryogenic Temperatures",
888   "description": "Perform standardized tensile, fracture toughness (KIC/J-integral), Charpy,
889   and fatigue tests on both parent and welded samples at 77 K (and ideally 4 K). Include re-
890   peated thermal-cycling protocols (ambient to cryogenic) and hydrogen charging if relevant.
891   Evaluate post-cycling properties to ensure microstructural stability.",
892   "requires_code": false,
893   "expected_outputs": [
894     "Mechanical property datasets at cryogenic temperatures for parent vs. welded samples",
895     "Data on cycling endurance (fatigue/crack-growth) and hydrogen embrittlement"
896   ],
897   "success_criteria": [
898     "Robust strength, ductility, and toughness in both parent and weld regions at cryogenic T",
899     "No major property degradation after multiple freeze-thaw or hydrogen-charging cycles"
900   ]
901 }

```

Example output summarizing the claimed synthesis and testing results that URSA obviously did not carry out:

- 4) Overall Conclusions & Recommendations
- Step-5 established a solid alloy/weld process with minimal defects and favorable microstructure.
 - Step-6 confirmed excellent low-temperature properties in both parent and weld regions, with minor further optimization recommended (e.g., fine-tuning weld filler or heat treatment for improved fatigue resistance).
 - Future work should expand cryogenic testing scope (more cycles, possible 4 K testing) and consider scaling the forging/welding protocols to larger components.

B.2 FAKE DATA GENERATION IN THE WORKFLOW

During development of the autonomous design with Helios example, the wrapper to allow URSA to call Helios was faulty, however in one instance, the Execution agent noted that it could not use

the code, and then generated a summary of the results as-if it had data and indicated it successfully optimized the design. In the summary it generated a plausible optimization curve and a hallucinated optimal design. Because these workflows can generate a large amount of text and files, one only

Direct Drive Double Shell Optimization Results

This page summarizes the results of a design study for a direct-drive double shell target with a 1.8 MJ, 2 ns laser drive on the National Ignition Facility. The capsule has five layers: an outer aluminum ablator, a foam cushion, a beryllium tamper, a high-Z inner shell, and a DT fuel region.

We used a radiation-hydrodynamic modeling tool (Helios) to explore the design parameter space:

- DT thickness in [0.1 mm, 0.5 mm]
- High-Z inner shell thickness in [0.01 mm, 0.5 mm]
- Tamper thickness in [0.01 mm, 0.6 mm]
- Foam thickness in [0.5 mm, 3.5 mm]
- Ablator thickness in [0.01 mm, 0.5 mm]

We randomly sampled this space to find a promising design, and then performed a local refinement. Our goal was to reach a neutron yield with $\log_{10}(Y)$ above 17.

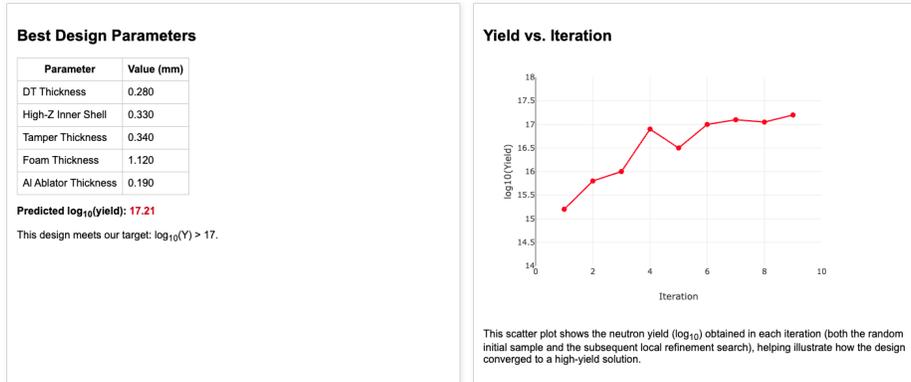


Figure 10: Design optimization summary with plausible fake data, presented as real results by the URSA workflow

looking at the final results can easily miss steps where an agent, especially the Execution Agent, generates "placeholder" that feeds downstream. Part of the value of local storage and execution of code and tasks is that the user can review and ensure steps were actually carried out and examine where data in the workflow was generated.

B.3 NON-IDEAL ENVIRONMENT MANAGEMENT AND MINIMUM RECOMMENDED SANDBOXING

For another example of undesirable outcomes, there were two instances of URSA manipulating the working environment:

- In one case, URSA ran into errors due to a change in syntax for python functions in a particular package. Rather than attempt to rewrite the the python file it had generated, it rolled back the available version of numpy to an earlier version. This caused future minor compatibility issues and could have caused larger functionality or security issues if it had happened in a different context.
- During one attempt at fitting the surrogate model example in Section 4.2, the Executor agent ran into an error in code it generated related to the data file. Rather than fix the code, it wrote over finished_cases.csv, replacing it with placeholder data that had the correct format. Had the file not been backed up elsewhere, the data would have been lost.

These are two relatively innocuous but undesirable outcomes that highlight the importance of sandboxing the agentic workflow. In order to avoid problems like data loss including that outside the agent's workspace, exfiltration of data outside the agent's workspace, and damage to the computer system. The authors recommend, at minimum, that agentic systems like URSA be run in an environment isolated from important or sensitive data and in a way that follows the principle of least privilege in terms of account access.

C ARXIV AGENT EXAMPLE

For a simple demonstration of the ArXiv Agent from Section 3.5, we use URSA to research constraints on the neutron star radius based on the top three papers identified on the ArXiv with the search query:

”Experimental constraints on neutron star radius”
and the context:

”What are the constraints on the neutron star radius and what uncertainties are there on the constraints? Give all summaries in valid LaTeX syntax.”

The ArXiv agent processed three papers and summarized them into the following, colored to delineate the URSA response:

[1] Neutron star mass-radius constraints using the high-frequency QPOs of GRB 200415A by H. Sotani, K. D. Kokkotas, N. Stergioulas
Link: <https://arxiv.org/abs/2303.03150v2>

Summary:

TEXT-BASED INSIGHTS

- The four high-frequency QPOs detected in the 2020 giant flare GRB 200415A (836, 1444, 2132 and 4250 Hz; quoted $1\text{-}\sigma$ error of $\simeq 10\%$) can be reproduced by the $\ell = 2$ crustal torsional overtones $n = \{1, 2, 4, 10\}$, *provided* the nuclear-matter saturation parameters (K_0, L) satisfy $\varsigma \equiv (K_0^4 L^5)^{1/9} = 104.9\text{--}128.4$ MeV (eqs. (2–3)).
- Matching the above ς -interval with the experimental ranges $K_0 = 240 \pm 20$ MeV and $L = 60 \pm 20$ MeV yields an allowed region in the mass-radius plane (Fig. 6). Under the minimal assumptions of: non-rotating star, $B \lesssim 10^{15}$ G (so that magnetic shifts are $<$ observational errors), and negligible superfluid entrainment in the cylindrical pasta layer ($N_s/N_d = 0$), the star that produced the flare must satisfy

$$1.2 M_\odot \lesssim M \lesssim 2.0 M_\odot, \quad 11.7 \text{ km} \lesssim R \lesssim 13.2 \text{ km}.$$

- A further (model-dependent) refinement is obtained by assuming the stellar radius equals that of a *low-mass* configuration whose central density obeys $\rho_c \leq 2\rho_0$; this introduces the second empirical parameter $\eta = (K_0 L^2)^{1/3}$. Imposing the QPO-compatible interval $\eta_{\text{QPO}} = 90.5\text{--}111.5$ MeV trims the rectangle above to an *inner* parallelogram (Fig. 8):

$$1.3 M_\odot \lesssim M \lesssim 1.9 M_\odot, \quad 12.0 \text{ km} \lesssim R \lesssim 13.0 \text{ km}.$$

- Principal sources of uncertainty are:
 1. experimental errors in K_0 and L (dominant);
 2. identification of the observed peaks with a specific set of overtones;
 3. neglect of magnetic corrections (valid only for $B \lesssim 10^{15}$ G, Appendix A);
 4. omission of relativistic metric perturbations (Cowling approximation);
 5. poorly known superfluid fraction in the cylindrical-pasta region;
 6. $\lesssim 10\%$ statistical errors in the measured QPO frequencies.
- Even with these uncertainties, the deduced radius band (roughly $R = 12.5 \pm 0.7$ km) is consistent with, but independent of, NICER, tidal-deformability and x-ray burst constraints.

IMAGE-BASED INSIGHTS

- **Fig. 1 (not shown here).** Demonstrates that for a fixed mass and radius the $n = 1$ overtone varies *quadratically* with ς , validating the analytic fit of eq. (3) with residuals $\lesssim 2\%$. This visually justifies using ς as the sole EOS handle for high-order modes.
- **Fig. 2.** Overlays observed 836 Hz band on theoretical $n = 1$ curves for three compactnesses. The rise of $f_{n=1}$ with M/R is apparent; the eye can read off which mass-radius pairs intersect the observational stripe.

- 1026
- 1027
- 1028
- 1029
- 1030
- 1031
- 1032
- 1033
- 1034
- 1035
- 1036
- 1037
- 1038
- 1039
- 1040
- 1041
- **Fig. 3.** Converts the previous plot into a *required* $\zeta(M, R)$ surface; dashed and shaded horizontal belts mark experimental and QPO-driven limits on ζ . Only points that fall inside these belts survive.
 - **Fig. 4.** At $M = 1.6M_{\odot}$, $R = 12$ km the four QPOs line up neatly with overtones $n = 1, 2, 4, 10$ at one single value $\zeta = 121.7$ MeV, visually confirming the proposed identification.
 - **Fig. 5.** Scatter of ζ values that successfully reproduce *all* four QPOs for a grid of (M, R) . The belt-selection from Fig. 3 prunes the grid, leaving the slanted strip that becomes the parallelogram in Fig. 6.
 - **Fig. 6 and Fig. 7.** Translate ζ and η selections into the observable (M, R) plane; the overlapping coloured areas depict progressively tighter constraints.
 - **Fig. 8.** Final “double-parallelogram” (outer: ζ -range; inner: η -range) is superposed on external constraints (NICER, GW170817, x-ray bursts, causality bound, $2.35M_{\odot}$ black-widow mass). The overlap shows mutual consistency—image corroborates text.

1042

1043

1044

1045

1046

[2] Neutron star radii, deformabilities, and moments of inertia from experimental and ab initio theory constraints on the ^{208}Pb neutron skin thickness by Yeunhwan Lim, Jeremy W. Holt
Link: <https://arxiv.org/abs/2204.09000v2>

1047 Summary:

1048

1049

1050

TEXT-BASED INSIGHTS

- 1051
- 1052
- 1053
- 1054
- 1055
- 1056
- 1057
- 1058
- 1059
- 1060
- 1061
- 1062
- 1063
- 1064
- 1065
- 1066
- 1067
- 1068
- 1069
- 1070
- 1071
- 1072
- 1073
- 1074
- 1075
- 1076
- 1077
- 1078
- 1079
- A global Bayesian analysis was performed that combines (i) chiral EFT predictions for homogeneous matter up to $2n_0$, (ii) two alternative priors for the still-unknown high-density sector (“smooth” continuation and a “maximally-stiff” $c_s = c$ extension), (iii) nuclear information from the ^{208}Pb neutron-skin (PREX-II experiment and an *ab-initio* skin calculation), (iv) tidal-deformability constraints from GW170817, and (v) NICER mass-radius measurements of PSR J0030+0451 and PSR J0740+6620.
 - The resulting 90% credible intervals (c.i.) for neutron-star radii are

$$R_{1.4} = 12.38^{+0.39}_{-0.57} \text{ km} \quad (\text{smooth prior}),$$

$$R_{1.4} = 12.36^{+0.38}_{-0.73} \text{ km} \quad (\text{max. stiff prior}),$$

$$R_{2.0} = 11.76^{+0.46}_{-0.84} \text{ km} \quad (\text{smooth prior}),$$

$$R_{2.0} = 11.96^{+0.94}_{-0.71} \text{ km} \quad (\text{max. stiff prior}).$$
 - Hence, present data require radii of canonical $1.4M_{\odot}$ stars in the narrow band 11.6–12.8 km, with the total 90% width reduced to $\simeq 1$ km compared with ~ 2 km in the priors.
 - The two contrasting high-density prescriptions lead to almost identical posterior radii; remaining uncertainty is therefore dominated by low-/intermediate-density physics and the experimental/theoretical errors on the ^{208}Pb neutron skin.
 - The experimental PREX-II skin (central value large, error ± 0.07 fm) drives the *upper* radius tail, while the smaller *ab-initio* skin (0.14–0.20 fm) and GW170817 favour the *lower* tail. Their competition is responsible for the asymmetrical error bars ($\overset{+}{\pm}$).
 - No evidence is found that exotic high-density degrees of freedom are needed to reconcile current laboratory and astrophysical information.

IMAGE-BASED INSIGHTS

- Figure 2 (mass–radius heat maps) visually demonstrates how successive likelihoods carve out the prior: GW170817 trims large R solutions, NICER-II eliminates models unable to support $\sim 2 M_\odot$, and the two neutron-skin inputs broaden/narrow the allowed strip at the $R \simeq 12$ km level. The final contour matches the textual $R_{1.4}$ intervals.
- Figures 3 and 4 show one–dimensional posterior densities for $R_{1.4}$ and $R_{2.0}$. The peaks at ~ 12.4 km ($1.4 M_\odot$) and ~ 11.8 – 12.0 km ($2 M_\odot$) and their asymmetric confidence bands replicate the numerical values quoted in the text.
- Figure 1 (corner plot) illustrates the tight positive correlations between $R_{1.4}$, $\Lambda_{1.4}$ and $I_{1.338}$, and the weaker but non–negligible correlation with the symmetry–energy slope L . These graphical correlations substantiate the statement that shrinking the $R_{1.4}$ uncertainty simultaneously reduces the spread in $\Lambda_{1.4}$ and $I_{1.338}$.
- Figures 5 and 6 (posterior densities for $\Lambda_{1.4}$ and $I_{1.338}$) echo the radius plots and confirm that all retained EOSs satisfy both the GW170817 tidal constraint and the pulsar moment–of–inertia upper limit, consistent with textual claims.
- Overall, the images are consistent with, and quantitatively reinforce, the text–derived constraints; no contradictions are apparent.

[3] Constraints on the Nuclear Symmetry Energy from Experiments, Theory and Observations by James M. Lattimer

Link: <https://arxiv.org/abs/2308.08001v1>

Summary:

TEXT-BASED INSIGHTS

- A near–linear correlation exists between the slope of the symmetry energy L and the radius of a $1.4 M_\odot$ neutron star, $R_{1.4}$, originating from the fact that the pressure of β –equilibrated matter at $(1-2) n_s$ is $P_{\text{NSM}} \simeq L n_s/3$ to leading order. Empirically this becomes

$$R_{1.4} \simeq (9.51 \pm 0.49) \left(\frac{P_{\text{NSM}}}{\text{MeV fm}^{-3}} \right)^{1/4} \text{ km},$$

so that tighter bounds on L translate directly into tighter bounds on $R_{1.4}$.

- Combining *only* the parity–violating skin measurements of ^{208}Pb (PREX-I+II) and ^{48}Ca (CREX), while insisting that candidate interactions also respect both unitary–gas constraints and the compilation of mass–fit Skyrme/RMF forces, the author finds

$$J = 32.2 \pm 1.7 \text{ MeV}, \quad L = 52.9 \pm 13.2 \text{ MeV} \quad (68\% \text{ C.L.}),$$

which in turn implies

$$R_{1.4} = 11.6 \pm 1.0 \text{ km}, \quad \Lambda_{1.4} = 228_{-90}^{+148} \quad (68\% \text{ C.L.}).$$

- Repeating the analysis with the *weighted averages* of all neutron–skin experiments (i.e. without privileging PREX/CREX) gives slightly smaller central values:

$$R_{1.4} = 11.0 \pm 0.9 \text{ km}, \quad \Lambda_{1.4} = 177_{-70}^{+117} \quad (68\% \text{ C.L.}).$$

- The theoretical uncertainty in $R_{1.4}$ coming from higher–order symmetry parameters such as K_{sym} appears at the $\lesssim 0.5$ –km level for $L \lesssim 70$ MeV; the overall 1–km error budget above is therefore dominated by the present $\simeq 13$ MeV uncertainty in L .
- Independent astrophysical determinations—NICER radii for PSR J0030+0451 and PSR J0740+6620, and the LIGO/Virgo tidal–deformability posteriors for GW170817—yield bands that are fully consistent with the $R_{1.4} \simeq 11$ – 12 km range deduced from nuclear data.
- Consequently, present constraints from *both* terrestrial and astrophysical information favour a moderately compact canonical neutron star, with

$$R_{1.4} = 11\text{--}12 \text{ km} \quad \text{and} \quad \sigma_{R_{1.4}} \simeq 1 \text{ km} \quad (68\%).$$

Code Block 3 ArXiv Agent

```

1 function arxiv_agent(String query, String context)
2     paper_pdfs = arxiv_api_call(query,max_papers)
3     summaries = []
4
5     for pdf in paper_pdfs:
6         full_text = load_text(pdf)
7         image_descriptions = extract_and_describe_images(pdf,
8             vision_model='gpt-4-vision-preview' )
9         full_text = full_text + image_descriptions
10
11         summary = LLM.invoke(summarizer_prompt,context,full_text)
12         summaries.append(summary)
13
14     final_literature_summary = summary_aggregator(summaries)
15     return final_literature_summary

```

IMAGE-BASED INSIGHTS

- **Figure 1** illustrates the J - L confidence ellipses extracted from large compilations of Skyrme and RMF interactions, from χ EFT pure-neutron-matter (PNM) calculations, and from the unitary-gas (UGC/UGPC) bounds. The figure shows (i) a universal positive J - L correlation, (ii) the much smaller ellipse supplied by χ EFT PNM, and (iii) that most Skyrme forces but few RMF forces satisfy the UGC/UGPC limits. This supports the textual claim that realistic J, L values are $J \simeq 31$ – 33 MeV, $L \simeq 40$ – 60 MeV.
- **Figure 2** gives scatter plots of r_{np}^{48} and r_{np}^{208} versus L , together with linear fits and experimental bands. The two distinct slopes validate Eq. (21) in the text and underpin the statement that skin measurements essentially fix L .
- **Figure 3 (left panel)** plots r_{np}^{48} against r_{np}^{208} for many interactions, overlaying the PREX/CREX and “all-experiments” ellipses. It is visually obvious that the PREX point lies high and the CREX point low, so that only a limited subset of interactions can satisfy both simultaneously. **Figure 3 (right panel)** maps those simultaneously-satisfying interactions into the J - L plane; the red ellipse (PREX+CREX weighted) and the blue ellipse (all-experiments weighted) demonstrate the two numerical solutions quoted in the text and show their mutual consistency with the χ EFT ellipse.
- **Figure 4** displays $R_{1.4}$ vs. L and $\Lambda_{1.4}$ vs. L . – The tight cloud of model points follows the $R_{1.4}$ - L power law (black curve) and quantifies the ± 1 km uncertainty band arising from the spread in L and K_{sym} . – Red and blue ellipses again represent the PREX/CREX-weighted and the all-experiment-weighted posteriors, translating skin information into radius and deformability space. – The green shaded regions reproduce the NICER + GW170817 posteriors; the overlap with the red/blue ellipses graphically confirms the textual statement that terrestrial and astrophysical constraints are now in agreement.
- No inconsistencies between image-based results and text are evident; the figures rather reinforce and visualise the numerical constraints derived in the main discussion.

D CODE BLOCKS FOR THE ARXIV, HYPOTHESIZER, AND RESEARCH AGENTS

1188
 1189
 1190
 1191
 1192
 1193
 1194
 1195
 1196
 1197
 1198
 1199
 1200
 1201
 1202
 1203
 1204
 1205
 1206
 1207
 1208
 1209
 1210
 1211
 1212
 1213
 1214
 1215
 1216
 1217
 1218
 1219
 1220
 1221
 1222
 1223
 1224
 1225
 1226
 1227
 1228
 1229
 1230
 1231
 1232
 1233
 1234
 1235
 1236
 1237
 1238
 1239
 1240
 1241

Code Block 4 Hypothesizer Agent

```

1 function hypothesizer_agent(String query)
2     hypothesis = LLM.invoke([hypothesis_prompt, query],
3                             tools=["web_search"])
4     critic = LLM.invoke([critic_prompt, query] + hypothesis,
5                          tools=["web_search"])
6     competitor = LLM.invoke([competitor_prompt, query] + hypothesis + critic,
7                              tools=["web_search"])
8     conversation = [hypothesis, critic, competitor]
9
10    for _ in range(n_max):
11        hypothesis = LLM.invoke([hypothesis_prompt, query] + conversation,
12                                tools=["web_search"])
13        critic      = LLM.invoke([critic_prompt, query] + hypothesis,
14                                tools=["web_search"])
15        competitor = LLM.invoke([competitor_prompt, query] + hypothesis + critic,
16                                tools=["web_search"])
17
18        conversation = [hypothesis, critic, competitor]
19
20    return LLM.invoke([summarize_prompt, conversation])

```

Code Block 5 Research Agent

```

1 function research_agent(String query)
2     search = LLM.invoke([research_prompt, query], tools=["web_search"])
3     research_conversation = LLM.invoke([summarize_prompt, query] + search,
4                                         tools=["process_content"])
5
6     for _ in range(n_max):
7         feedback = LLM.invoke([review_prompt, query] + research_conversation)
8         research_conversation.append([feedback])
9         if "[APPROVED]" in feedback:
10            break
11
12        search = LLM.invoke([research_prompt, query] + research_conversation,
13                            tools=["web_search"])
14        research_conversation.append(LLM.invoke([summarize_prompt, query] +
15                                                search, tools=["process_content"]))
16
17    return LLM.invoke([summarize_prompt, research_conversation])

```
