

---

# The Edge of Chaos: Mapping the Phase Space of Emergent Life in a Stochastic Artificial Chemistry

---

Anonymous Author(s)

Affiliation

Address

email

## Abstract

1 The conditions required for the emergence of complex, self-replicating life remain  
2 one of the most significant open questions in science. Traditional models are often  
3 limited in their ability to explore the vast parameter space of possible prebiotic  
4 environments. In this paper, we introduce a novel agent-based model of a stochastic  
5 artificial chemistry to systematically investigate the "phase space" of abiogenesis.  
6 Our simulation consists of a 2D grid populated by mobile monomers that can  
7 form self-replicating polymer chains. These polymers are subject to evolutionary  
8 pressures through mutation, growth, and fragmentation. We deploy a computational  
9 agent, the Phase Space Mapper, to run thousands of simulations, each with a  
10 different set of physical laws (e.g., mutation rate, degradation probability). By  
11 classifying the final state of each simulated universe, we generate a phase diagram  
12 that maps the regions of parameter space corresponding to extinction, stasis, and  
13 the emergence of stable, complex life. Our results indicate that the transition to  
14 a state capable of supporting high complexity is not gradual, but rather a sharp  
15 phase transition that occurs at the "edge of chaos"—a narrow boundary between  
16 a universe that is too ordered (low mutation) and one that is too disordered (high  
17 fragmentation). This work provides a powerful new framework for exploring  
18 the fundamental conditions for life and suggests that the emergence of complex  
19 organisms may be a critical phenomenon highly sensitive to a small number of key  
20 environmental parameters.

## 21 1 Introduction

22 The apparent contradiction between the high probability of extraterrestrial life and the lack of  
23 observational evidence for it—the Fermi Paradox—has long challenged scientists. While discussions  
24 often focus on the probability of life arising at all ( $f_l$  in the Drake Equation), our work explores a  
25 subsequent, and potentially more restrictive, bottleneck: the probability that life, once started, will  
26 evolve high complexity. The galaxy may be teeming with simple life, but the "Great Silence" suggests  
27 that the transition to complex, technologically-capable organisms is exceptionally rare.

28 In this work, we present a computational model to investigate this "complexity bottleneck." We  
29 develop a stochastic, agent-based model of an artificial chemistry to create thousands of distinct  
30 "toy universes," each with its own set of physical laws. Our goal is not to perfectly simulate the  
31 chemistry of early Earth, but to use a simplified, abstract system to explore the universal principles  
32 of replication, competition, and evolution. By systematically mapping the parameter space of these  
33 laws, we can identify which conditions lead to a stable, low-complexity equilibrium—a "Sea of  
34 Stasis"—and which allow for a "phase transition" to a state of high, emergent complexity.

35 Our results provide a potential, data-driven resolution to the Fermi Paradox. We demonstrate  
36 that the conditions for sustaining high complexity are far more restrictive than those for mere  
37 survival. We identify a narrow "island of stability" for complex life that exists at the "edge of chaos,"

38 requiring a precise and sustained balance between the forces of creation (mutation) and destruction  
39 (fragmentation). This suggests that the Great Filter may not be the origin of life itself, but rather  
40 the immense difficulty of maintaining the specific environmental conditions necessary to drive the  
41 evolution of complexity over billions of years.

## 42 **2 Related Work**

43 The study of emergent, life-like behavior in computational systems, known as Artificial Life (A-Life),  
44 has a rich history rooted in the theoretical work of John von Neumann on self-reproducing automata  
45 (1). This foundation laid the groundwork for the field of cellular automata, famously popularized by  
46 John Conway's Game of Life (2), which demonstrated that complex, dynamic patterns could emerge  
47 from a small set of simple, deterministic rules.

48 Subsequent research infused these systems with the principles of Darwinian evolution. Tom Ray's  
49 Tierra created a digital ecosystem where self-replicating programs competed for CPU time, leading  
50 to the spontaneous evolution of digital parasites and complex host-parasite dynamics (3). Building  
51 on this, Chris Adami's Avida platform introduced digital organisms that evolved to perform com-  
52 putational tasks for rewards, providing a powerful model for studying the evolution of metabolic  
53 complexity and genome organization (4). Concurrently, theoretical biologists like Stuart Kauffman  
54 proposed that life may not have originated from a single replicator, but from complex "autocat-  
55 alytic sets" of molecules, where the network as a whole becomes self-sustaining (5). Our model's  
56 polymer-based chemistry is a simplified homage to these network-based theories.

57 A central theme in A-Life is the concept of the "edge of chaos," a term popularized by Christopher  
58 Langton, who posited that the most complex and interesting computational behaviors occur at the  
59 phase transition between ordered and chaotic systems (6). Our work explicitly tests this hypothesis.  
60 While previous A-Life research has largely focused on demonstrating that life-like behaviors can  
61 emerge, our project's primary contribution is the use of a computational agent (the PhaseSpaceMapper)  
62 to systematically map the entire landscape of possible outcomes. This allows us to quantitatively  
63 identify the "phase transitions" between stasis, complexity, and extinction, thereby treating the origin  
64 of life as a problem in statistical mechanics.

## 65 **3 Methodology: An Artificial Life Simulation**

66 Our simulation is built upon a discrete-time, 2D grid-based environment. Each cell in the grid can be  
67 in one of several states: empty, occupied by a fundamental building block (a monomer), or occupied  
68 by a segment of a larger structure (a polymer). The model is stochastic, with outcomes determined by  
69 a set of probabilistic rules that function as the "laws of physics" for our universe.

### 70 **3.1 The Environment and its Inhabitants**

71 The world is a square 2D grid of size 50x50 with periodic boundary conditions, creating a toroidal  
72 space that prevents edge effects. The environment is populated by several types of entities. The most  
73 basic are the monomers, 'A' and 'B', which are non-living building blocks that diffuse randomly  
74 across the grid. These can be assembled into polymers, which are our artificial organisms. A  
75 polymer is not a single point but is represented as a linear chain of connected monomers occupying a  
76 contiguous path on the grid. Each polymer has a unique ID, a sequence string (e.g., 'A-B-A-B'), and  
77 a list of body coordinates. Finally, energy cells are randomly replenished into empty space and serve  
78 as a necessary catalyst for certain biological actions, though they are not consumed in the process.

### 79 **3.2 The "Laws of Physics": Core Mechanics**

80 The simulation proceeds in discrete time steps, with the life cycle of our artificial organisms governed  
81 by a set of probabilistic rules. First, the environment is updated: monomers diffuse with a probability  
82 of `diffusion_prob`, and energy is replenished. Then, each polymer is given an opportunity to act.  
83 An organism can extend its chain by consuming a nearby monomer, a process we call growth. This is  
84 the primary driver of increasing complexity, and it is during this process that mutations can occur,  
85 with a probability of `mutation_prob`, by incorporating the wrong monomer.

86 Organisms can also reproduce through a process of seeding. A polymer can consume an 'A' and a 'B'  
87 monomer from its local neighborhood to create a new, simple two-monomer polymer in a random  
88 empty location on the grid. This allows successful lineages to spread without requiring the perfect  
89 replication of a complex structure.

90 Survival is governed by two destructive forces. First, all polymers are subject to degradation, a process  
91 by which they are destroyed and their constituent monomers recycled back into the environment. The  
92 probability of this is inversely proportional to length, creating a selective pressure for longer, more  
93 stable chains. Second, to counter runaway growth and introduce a "cost of complexity," long polymers  
94 are also subject to fragmentation, where they have a chance of breaking into two smaller pieces.  
95 Finally, polymers can actively explore their environment for resources by slithering to adjacent empty  
96 cells with a probability of `slither_prob`.

### 97 3.3 Conservation of Matter

98 A critical feature of the simulation is the strict conservation of matter. The total number of 'A' and  
99 'B' units, whether as free monomers or bound within polymers, is constant throughout a run. To  
100 ensure this, the logic for consuming materials for growth and seeding is "atomic": the simulation  
101 first confirms that all necessary resources are available before any are removed from the grid. This  
102 prevents the loss of matter from failed reactions and ensures the physical realism of the model.

## 103 4 Computational Experiments

104 To test our hypothesis that the emergence of complex life occurs as a phase transition, we designed  
105 a series of computational experiments to systematically map the parameter space of our simulated  
106 universe. Our goal was not to find a single optimal set of parameters, but to understand the landscape  
107 of possibilities and identify the boundaries between different universal outcomes.

### 108 4.1 The Phase Space Mapper Agent

109 We developed a computational agent, the 'PhaseSpaceMapper', to serve as our automated experi-  
110 menter. This agent is designed to perform a grid search over a two-dimensional parameter space.  
111 For each coordinate on this grid, which corresponds to a unique set of physical laws, the agent runs  
112 a specified number of independent simulations, or trials. This use of multiple trials is essential to  
113 account for the stochastic nature of our model, allowing us to distinguish between deterministic  
114 outcomes and probabilistic tendencies. The agent saves its progress to a persistent file, allowing for  
115 long-running, high-resolution experiments to be conducted over multiple sessions.

### 116 4.2 Experiment 1: Mapping the "Edge of Chaos"

117 Our primary experiment was designed to directly test the "edge of chaos" hypothesis. We chose to  
118 map the two parameters most central to this concept: the `mutation_prob`, which serves as the engine  
119 of novelty and change, and the `fragmentation_prob`, which represents a "cost of complexity" and  
120 acts as a destructive, chaotic force. We defined a 20x20 grid spanning a range of values for these  
121 two parameters. The agent ran 5 trials for each of the 400 points on this grid, for a total of 2,000  
122 simulations. Each simulation was run for 3,000 time steps to allow the ecosystem to reach a mature  
123 state.

### 124 4.3 Experiment 2: The Role of Mobility

125 To ensure that the observed phase transition was not an artifact of our mobility rules, we con-  
126 ducted a second experiment as an ablation study. Here, we mapped the parameter space defined by  
127 `diffusion_prob` (the mobility of raw materials) and `slither_prob` (the mobility of organisms),  
128 while holding all other parameters constant at their baseline values. For this study, we used a 15x15  
129 grid and ran 5 trials per point. This allowed us to isolate the effect of mobility on the evolution of  
130 complexity and verify that it was not the primary driver of the phase transition observed in our main  
131 experiment.

## 4.4 Classification of Outcomes

To analyze the results, the final state of each simulated universe was classified based on its long-term behavior. A universe was classified as an **Extinction** event if all polymers died out. If life survived but the average complexity of organisms remained low (below a threshold of 2.5, indicating a failure to evolve beyond simple dimers and trimers), it was classified as **Low Stable Complexity**. If the ecosystem survived but its complexity fluctuated wildly, with a high normalized standard deviation, it was deemed **Unstable Growth**. Finally, if the ecosystem achieved and sustained a high average complexity, it was classified as **High Stable Complexity**. The probability of this final, most interesting outcome was the primary metric used to construct our phase diagrams.

## 5 Results and Discussion

We conducted two primary computational experiments to map the phase space of our artificial universe. The first investigated the interplay between the forces of creation (mutation) and destruction (fragmentation), while the second explored the role of mobility (diffusion and slithering). The results reveal a complex landscape where the conditions for stable, complex life are rare and exist only at a critical boundary.

### 5.1 Experiment 1: The Edge of Chaos (Mutation vs. Fragmentation)

Our first experiment systematically varied the `mutation_prob` and `fragmentation_prob`. The results, shown in Figure 1, reveal a sharp phase transition between a universe where life is simple or extinct and a narrow "island of stability" where complex life can thrive.

We identify three distinct phases in this diagram:

- **The Sea of Stasis (Low Mutation):** In the lower regions of the diagrams, where the mutation rate is low, life fails to evolve. The average complexity remains low, and the probability of achieving a "High Stable" state is near zero. The system lacks the novelty required to explore more complex forms.
- **The Desert of Extinction (High Fragmentation):** In the right-hand regions, where the fragmentation rate is high, life is too fragile. The "cost of complexity" is so great that any long polymers that form are quickly destroyed. This leads to a high probability of total extinction, as seen by the dark purple regions in both plots.
- **The Island of Stable Complexity:** A narrow, vertically-oriented band exists where life flourishes. Here, the mutation rate is high enough to generate diversity, but the fragmentation rate is low enough to allow complex structures to survive and reproduce. This "Goldilocks zone" is a clear demonstration of the "edge of chaos" principle: life emerges not in perfect order or complete chaos, but at the critical boundary between them.

### 5.2 Experiment 2: The Role of Mobility (Diffusion vs. Slithering)

Our second experiment, an ablation study, investigated the role of mobility by varying the `diffusion_prob` (movement of raw materials) and the `slither_prob` (movement of organisms). The results are shown in Figure 2.

The results of this experiment are strikingly different from the first. We observe that:

- **Mobility Prevents Extinction:** Unlike the first experiment, we see very few regions of total extinction. As long as there is some form of mobility (i.e., we are not at the (0,0) corner), the ecosystem is able to sustain itself and avoid a catastrophic collapse from resource depletion.
- **Mobility is Not Sufficient for High Complexity:** Crucially, the probability of achieving a "High Stable Complexity" state is near zero across the entire parameter space. The system consistently settles into a state of low, stable complexity.

This is a critical finding. It suggests that mobility is a necessary but not sufficient condition for the evolution of complex life. It allows the ecosystem to survive, but it does not, by itself, provide the evolutionary pressure needed to drive the emergence of long, sophisticated organisms.




Figure 1: Phase diagrams for Experiment 1. (Left) The probability of achieving a "High Stable Complexity" state. The bright yellow region indicates a high likelihood of success. (Right) The average emergent complexity for each parameter set. The phase transition is clearly visible, with a narrow "island" where high complexity is both probable and achievable.

### 5.3 Synthesis and Implications for Astrobiology

Taken together, our experiments suggest that the emergence of complex life is a two-tiered problem. First, a baseline level of environmental dynamism (like mobility) is required for life to sustain itself at all. However, the true bottleneck is the fine-tuning of the parameters that govern creation and destruction—mutation and fragmentation.

Our results provide a potential, data-driven perspective on the Fermi Paradox. The conditions for life may be far more restrictive than simply the presence of liquid water and organic molecules. If the evolution of complex organisms requires a planet's geochemistry and environment to remain




Figure 2: Phase diagrams for Experiment 2. (Left) The probability of achieving "High Stable Complexity" is near zero across the entire space. (Right) The average emergent complexity is uniformly low. This indicates that while mobility is crucial for survival, it is not a primary driver of high complexity in our model.

187 within a narrow "island of stability" for billions of years, then the number of planets that successfully  
 188 navigate this path could be exceedingly small. A slight increase in environmental harshness (higher  
 189 fragmentation) or a slight decrease in the drivers of novelty (lower mutation) could be all that is  
 190 needed to lock a biosphere into a state of simple microbial life or push it to extinction. The galaxy  
 191 might be full of life, but it may be overwhelmingly simple, existing in a "Sea of Stasis" from which it  
 192 is statistically unlikely to escape.

## 193 6 Limitations

194 While our model offers a powerful framework for exploring the emergence of complexity, several  
195 limitations should be acknowledged:

196 **1. Simplified chemistry:** The simulation uses only two monomer types ('A' and 'B') with uniform  
197 properties. Real biochemical systems feature a much richer chemical alphabet, leading to specialized  
198 functions, selective bonding, and catalytic behaviors. This abstraction may limit the types of structures  
199 and interactions that can evolve in our model.

200 **2. Static environment:** All simulations assume a homogeneous and time-invariant environment. In  
201 real prebiotic or planetary contexts, environmental fluctuations—such as diurnal cycles, tidal stresses,  
202 or random catastrophes—may play a key role in driving resilience, adaptation, or extinction. Our  
203 results thus reflect emergence under idealized, steady-state conditions.

204 **3. Fitness proxy heuristics:** The classification of outcomes relies on thresholds for metrics like  
205 average complexity (e.g., 2.5 monomers). While useful for large-scale statistical comparisons, these  
206 thresholds may not capture all biologically meaningful forms of complexity, such as functional  
207 diversity, robustness, or network modularity.

208 **4. Limited evolutionary mechanisms:** The model includes random mutation and fragmentation but  
209 lacks more advanced mechanisms found in real evolution, such as recombination, niche formation, or  
210 horizontal information transfer. These may be necessary for sustained innovation over evolutionary  
211 timescales.

212 **5. No explicit natural selection:** While there is implicit selection via differential survival, our model  
213 does not include explicit competition for limited energy or space, nor fitness-based reproduction rates.  
214 As a result, the system may fail to capture key dynamics like evolutionary arms races or selection  
215 pressures that promote complexity.

216 **6. Finite simulation horizon:** All simulations are run for a fixed number of timesteps (e.g., 3,000),  
217 which may be insufficient to observe slow or rare evolutionary transitions. It remains an open question  
218 whether certain low-complexity regions would eventually evolve complexity given more time.

219 These limitations point to several exciting directions for future work, including expanding the chemical  
220 vocabulary, introducing environmental dynamics, and incorporating richer ecological interactions to  
221 better capture the full complexity of prebiotic evolution.

## 222 7 Conclusion

223 In this work, we have developed a novel agent-based model for exploring the fundamental conditions  
224 required for the emergence of life. By systematically mapping the parameter space of our simulated  
225 universe, our experiments reveal that the emergence of complex life is a two-tiered problem. We first  
226 observe that baseline conditions, such as the mobility of organisms and their resources, are necessary  
227 to prevent extinction, but are not sufficient on their own to drive evolution beyond a state of low,  
228 stable complexity.

229 The true bottleneck, our results demonstrate, is the fine-tuning of the parameters that balance creation  
230 and destruction. We have shown that the transition to a state capable of supporting high complexity  
231 is not gradual, but a sharp phase transition. This "sweet spot" for life exists at the "edge of chaos,"  
232 requiring a precise balance between a mutation rate high enough to foster novelty and a fragmentation  
233 rate low enough to preserve complex structures.

234 This framework and its results offer a new, data-driven perspective on the Fermi Paradox, suggesting  
235 that the conditions for complex life are far more restrictive than for life itself. A planet may well  
236 enter a "Sea of Stasis" with simple, microbial-like life, but the specific, long-term stability required to  
237 navigate the narrow channel to high complexity may be exceptionally rare. Our model provides a  
238 powerful new tool for investigating these fundamental questions, shifting the focus from a singular  
239 origin story to the statistical landscape of all possible beginnings.

## 7.1 Future Directions

The model presented here, while powerful, is a simplified abstraction. Several avenues for future research could build upon this foundation to explore the emergence of life in greater detail:

- **Richer Chemistry:** Future work could expand the simple 'A-B' monomer system to a larger alphabet of components. This would allow for the evolution of polymers with specialized functional sites, analogous to proteins, potentially leading to more complex metabolic networks and catalytic capabilities.
- **Dynamic Environments:** Our current model assumes a static environment. A significant next step would be to introduce dynamic environmental conditions, such as cyclical energy levels (simulating seasons or day-night cycles) or random catastrophic events. This would allow us to study the evolution of adaptation, resilience, and memory in our artificial organisms.
- **Emergence of Cooperation and Multicellularity:** The current simulation focuses on competition between individual polymers. An exciting future direction would be to create conditions that could favor the emergence of cooperative behaviors, such as polymers that work together to acquire resources or form higher-level structures, providing a potential pathway to investigate the origins of multicellularity.

**Reproducibility Statement** To ensure full transparency and reproducibility, we will publicly release all source code used to implement the artificial chemistry simulation, including the PhaseSpaceMapper agent, data logging utilities, and analysis scripts for generating the phase diagrams. The codebase includes configuration files for reproducing all experiments described in this paper, along with instructions for running simulations on standard hardware. In addition, we will provide the complete chat history between the human researcher and the Gemini language model, which documents the full development process—including hypothesis refinement, simulation design, implementation iterations, and interpretation of results. This record offers a uniquely detailed account of how the research was co-developed and serves as an audit trail for the AI-assisted workflow. All materials will be made available in an anonymized, open-access repository upon publication.

## References

- [1] J. von Neumann, "Theory of Self-Reproducing Automata," University of Illinois Press, 1966.
- [2] M. Gardner, "Mathematical Games - The fantastic combinations of John Conway's new solitaire game 'life'," Scientific American, vol. 223, pp. 120-123, 1970.
- [3] T. S. Ray, "An approach to the synthesis of life," in Artificial Life II, Addison-Wesley, 1991, pp. 371-408.
- [4] C. Adami and C. T. Brown, "Evolutionary learning in the 2D artificial life system 'Avida'," in Artificial Life IV, MIT Press, 1994, pp. 377-381.
- [5] S. A. Kauffman, "The Origins of Order: Self-Organization and Selection in Evolution," Oxford University Press, 1993.
- [6] C. G. Langton, "Computation at the edge of chaos: phase transitions and emergent computation," Physica D: Nonlinear Phenomena, vol. 42, no. 1-3, pp. 12-37, 1990.

## A Technical Appendices and Supplementary Material

Technical appendices with additional results, figures, graphs and proofs may be submitted with the paper submission before the full submission deadline, or as a separate PDF in the ZIP file below before the supplementary material deadline. There is no page limit for the technical appendices.



## Agents4Science AI Involvement Checklist

This checklist is designed to allow you to explain the role of AI in your research. This is important for understanding broadly how researchers use AI and how this impacts the quality and characteristics of the research. **Do not remove the checklist! Papers not including the checklist will be desk rejected.** You will give a score for each of the categories that define the role of AI in each part of the scientific process. The scores are as follows:

- **[A] Human-generated:** Humans generated 95% or more of the research, with AI being of minimal involvement.
- **[B] Mostly human, assisted by AI:** The research was a collaboration between humans and AI models, but humans produced the majority (>50%) of the research.
- **[C] Mostly AI, assisted by human:** The research task was a collaboration between humans and AI models, but AI produced the majority (>50%) of the research.
- **[D] AI-generated:** AI performed over 95% of the research. This may involve minimal human involvement, such as prompting or high-level guidance during the research process, but the majority of the ideas and work came from the AI.

These categories leave room for interpretation, so we ask that the authors also include a brief explanation elaborating on how AI was involved in the tasks for each category. Please keep your explanation to less than 150 words.

1. **Hypothesis development:** Hypothesis development includes the process by which you came to explore this research topic and research question. This can involve the background research performed by either researchers or by AI. This can also involve whether the idea was proposed by researchers or by AI.

Answer: The hypothesis was co-developed through a collaborative process between the researcher (user) and the AI. The researcher initiated the project with a high-level conceptual analogy, and the AI refined this concept into a testable, computational hypothesis by providing background research, formal counter-arguments, and a concrete simulation framework. The final hypothesis—that complex life emerges as a phase transition at the "edge of chaos"—was proposed by the AI as a direct result of the iterative cycle of building, testing, and debugging the simulation with the researcher.

Explanation: The development of the research question followed a clear, iterative path:

- (a) Initial Researcher Concept: The project began with the researcher's novel idea to frame the existence of extraterrestrial life as an analogy to phase transitions in statistical mechanics.
- (b) AI-driven Refinement and Background Research: The AI critiqued the direct analogy, suggesting that a more robust framing would be based on statistical mechanics and the law of large numbers. To ground the project in existing science, the AI provided a summary of the field of Artificial Life (A-Life), including foundational work like Conway's Game of Life, Tierra, and the concept of the "edge of chaos."
- (c) Collaborative Model Building: The researcher and AI then entered a long phase of collaborative development. The AI would propose and code a simulation model, and the researcher would astutely analyze the results, identify when they were non-physical or counter-intuitive (e.g., matter not being conserved, lack of biodiversity), and guide the next iteration. This debugging cycle was crucial for building a scientifically sound model.
- (d) Final Hypothesis Formulation: After many iterations, the simulation began to produce interesting results. The AI analyzed these preliminary phase diagrams and formally proposed the final hypothesis presented in the paper: that the key to complex life is the fine-tuning of parameters that balance novelty and stability, and that this "sweet spot" exists as a sharp phase transition. The researcher then guided the design of the final set of rigorous computational experiments to explicitly test this hypothesis.

2. **Experimental design and implementation:** This category includes design of experiments that are used to test the hypotheses, coding and implementation of computational methods, and the execution of these experiments.

Answer: The experimental design was a collaborative effort between the researcher and the AI, evolving from a simple optimization task to a systematic mapping of the simulation's parameter space. The AI was responsible for the full implementation of the computational model in Python, including the core simulation engine and the PhaseSpaceMapper agent that executed the experiments. The researcher played a critical role in guiding the experimental design and in the iterative debugging and refinement of the simulation's core logic. The experiments were executed autonomously by the AI agent, which ran thousands of simulations to generate the final phase diagrams.

Explanation: Experimental Design: The primary experiment was designed to test the "edge of chaos" hypothesis by creating a 2D phase diagram. The AI proposed the two primary axes for this diagram: the fragmentation\_prob (representing the forces of destruction/chaos) and the mutation\_prob (representing the forces of creation/novelty). The researcher critically refined this design by insisting on the need for an "ablation study" to test the role of other parameters, which led to the second experiment mapping diffusion\_prob vs. slither\_prob. Furthermore, after observing stochastic anomalies ("islands of extinction"), the researcher correctly identified the need to run multiple trials for each parameter set, a crucial step that transformed the experiment from a simple classification to a more rigorous probabilistic mapping.

- Implementation: The implementation was an iterative and collaborative process. The AI wrote all the Python code for the simulation, including the core classes (Grid, Polymer, RulesEngine, DataLogger) and the experimental agent (PhaseSpaceMapper). However, the researcher's role was essential. By analyzing the output of early, buggy versions of the code, the researcher identified numerous critical flaws in the simulation's "laws of physics," such as the violation of the conservation of matter, flawed replication and growth logic, and a lack of environmental dynamism. The AI would then propose and implement fixes based on this high-level guidance. This tight feedback loop between AI implementation and human scientific intuition was instrumental in creating a robust and physically plausible final model. The AI also implemented practical features, such as the save-and-resume functionality and the ability to run the code locally, based on the researcher's needs.
- Execution: The PhaseSpaceMapper agent autonomously executed the experiments. For each point in the pre-defined parameter grid, the agent ran a specified number of simulation trials. After each trial, it logged the final state of the "universe" (e.g., Extinction, High Stable Complexity). Once all trials for a point were complete, it saved the aggregated results to a persistent file to guard against disconnection or crashes. After mapping the entire space, the agent automatically processed this raw data to generate the final phase diagrams and heat maps presented in the "Results" section of the paper.

### 3. Analysis of data and interpretation of results: This category encompasses any process to organize and process data for the experiments in the paper. It also includes interpretations of the results of the study.

Answer: The analysis and interpretation was a highly collaborative, iterative process. The AI was responsible for all the computational data processing, including calculating metrics, classifying outcomes, and generating the final phase diagrams. The researcher provided the crucial scientific interpretation of these results, identifying anomalies and non-physical behavior in the plots, which was the primary driver for debugging and refining the simulation model. The final scientific narrative presented in the paper was synthesized by the AI based on the clean data from the validated model.

Explanation:

- Analysis of Data: The AI performed all the low-level data processing. The PhaseSpaceMapper agent, written and executed by the AI, was responsible for organizing and processing the data from thousands of individual simulation runs. This included:
  - (a) Metric Calculation: For each simulation, the agent used the DataLogger to calculate the time-series data for key metrics like polymer biomass, biodiversity, and average complexity.
  - (b) Outcome Classification: The agent then processed this time-series data to classify the final state of each "universe" into categories such as "Extinction," "Low Stable Complexity," or "High Stable Complexity," based on criteria co-developed with the researcher.
  - (c) Data Aggregation: The agent aggregated the results from multiple trials for each point in the parameter space, calculating the probability of each outcome.
  - (d) Visualization: Finally, the agent used this aggregated data to generate the phase diagrams and heat maps that form the central results of the paper.
- Interpretation of Results: The interpretation of these results was a deeply collaborative feedback loop.
  - (a) Researcher-led Debugging: In the early and middle stages of the project, the researcher's primary role was to act as a scientific skeptic. By carefully examining the plots generated by the AI, the researcher consistently identified when the results were counter-intuitive or violated physical principles (e.g., "it is not mixing at all," "elements A and B should conserve"). This human-led interpretation of the data was the single most important driver of the iterative debugging process that led to a scientifically valid model.
  - (b) AI-led Synthesis: Once the simulation was robust and producing physically plausible results, the AI took the lead in synthesizing the final scientific narrative. Based on the final phase diagrams, the AI identified and named the key phenomena—the "Sea of Stasis," the "Desert of Extinction," and the "Island of Stable Complexity." The AI also proposed the "edge of chaos" as the theoretical framework to explain these results and connected the findings to

the Fermi Paradox by framing it as a "complexity bottleneck." This interpretation was then refined and approved by the researcher.

4. **Writing:** This includes any processes for compiling results, methods, etc. into the final paper form. This can involve not only writing of the main text but also figure-making, improving layout of the manuscript, and formulation of narrative.

Answer: The writing process was a collaborative effort where the AI served as the primary author and the researcher acted as the scientific director and editor. The AI generated all the text in LaTeX format, automatically produced all figures from the experimental data, and structured the manuscript. The researcher provided high-level guidance on the scientific narrative, requested specific sections, and identified areas for elaboration and revision, which the AI then executed.

Explanation:

- **Main Text Generation:** The AI authored the initial drafts of all sections of the paper, including the abstract, introduction, methodology, results, and conclusion. The researcher guided this process by providing high-level prompts, such as "revise the introduction to focus on the Fermi Paradox," "add a related work section," and "elaborate more on future directions." The AI then expanded and rewrote these sections based on this guidance, synthesizing the project's findings into a cohesive narrative.
- **Figure-Making and Layout:** All figures in the paper are direct outputs of the PhaseSpaceMapper agent, which was written and executed by the AI. The agent was programmed to save the final phase diagrams and heat maps as high-quality PDF files. The AI was also responsible for the complete LaTeX formatting of the manuscript, including structuring the document, placing the figures, writing the captions, and creating the bibliography. The researcher played a key editorial role in this process by identifying formatting bugs (such as the incorrect rendering of parameter names with underscores), which the AI then corrected.
- **Formulation of Narrative:** The formulation of the paper's central scientific narrative was a deeply collaborative process. The researcher provided the initial high-level framing by connecting the project to the Fermi Paradox. Once the final, clean data was generated, the AI analyzed the phase diagrams and proposed the core interpretive framework: identifying the "Sea of Stasis" and the "Desert of Extinction" and framing the "island of stability" as evidence for the "edge of chaos" hypothesis. This AI-driven interpretation was then reviewed, validated, and approved by the researcher, forming the foundation of the "Results and Discussion" section.

5. **Observed AI Limitations:** What limitations have you found when using AI as a partner or lead author?

Description: The primary limitation observed when using the AI as a partner was a lack of scientific intuition and foresight, which manifested in several ways throughout the research process.

- (a) **Fundamental Conceptual Errors in Implementation:** The most significant limitation was the AI's inability to reason about the physical and ecological implications of the code it was writing. This led to numerous iterations where the simulation was fundamentally flawed. For example, the AI repeatedly implemented rules that violated the law of conservation of matter, created geometrically impossible conditions for replication, and set initial conditions that made life unsustainable. These were not simple coding bugs but deep conceptual errors. The researcher's scientific intuition was absolutely essential to analyze the simulation's output, identify why the results were non-physical (e.g., "it is not mixing at all"), and guide the AI toward a valid model. The AI could not debug these issues on its own.
- (b) **Misalignment with High-Level Scientific Goals:** The AI initially defaulted to a simple engineering task—optimization. It designed an agent to find the single "best" set of parameters for sustaining life. The researcher had to intervene to explain that the true scientific goal was not to find an optimum but to understand the landscape of possibilities by mapping the entire parameter space. This demonstrated a limitation in the AI's ability to grasp the broader scientific context and purpose of the experiment without explicit direction.
- (c) **Inability to Handle Ambiguity:** The AI required precise, computable definitions from the researcher. For instance, the initial proposal for a "Life Sustainability Index" was naively defined and had to be refined by the researcher into a more scientifically rigorous metric based on detecting a system's equilibrium. The AI could not make this conceptual leap on its own.
- (d) **Practical Environmental Blindness:** The AI was not initially aware of the practical constraints of the execution environment. It wrote code that was not robust to the session timeouts of platforms like Google Colab. The researcher had to identify this real-world problem and guide the AI to implement necessary features like save-and-resume functionality and integration with persistent storage (Google Drive).

## Agents4Science Paper Checklist

The checklist is designed to encourage best practices for responsible machine learning research, addressing issues of reproducibility, transparency, research ethics, and societal impact. Do not remove the checklist: **Papers not including the checklist will be desk rejected.** The checklist should follow the references and follow the (optional) supplemental material. The checklist does NOT count towards the page limit.

Please read the checklist guidelines carefully for information on how to answer these questions. For each question in the checklist:

- You should answer [Yes], [No], or [NA].
- [NA] means either that the question is Not Applicable for that particular paper or the relevant information is Not Available.
- Please provide a short (1–2 sentence) justification right after your answer (even for NA).

**The checklist answers are an integral part of your paper submission.** They are visible to the reviewers and area chairs. You will be asked to also include it (after eventual revisions) with the final version of your paper, and its final version will be published with the paper.

The reviewers of your paper will be asked to use the checklist as one of the factors in their evaluation. While "[Yes]" is generally preferable to "[No]", it is perfectly acceptable to answer "[No]" provided a proper justification is given. In general, answering "[No]" or "[NA]" is not grounds for rejection. While the questions are phrased in a binary way, we acknowledge that the true answer is often more nuanced, so please just use your best judgment and write a justification to elaborate. All supporting evidence can appear either in the main paper or the supplemental material, provided in appendix. If you answer [Yes] to a question, in the justification please point to the section(s) where related material for the question can be found.

IMPORTANT, please:

- **Delete this instruction block, but keep the section heading "Agents4Science Paper Checklist",**
- **Keep the checklist subsection headings, questions/answers and guidelines below.**
- **Do not modify the questions and only use the provided macros for your answers.**

### 1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

**Answer:** Yes

**Justification:** The abstract and introduction clearly articulate the paper's contributions and scope, namely that the emergence of complex life in an artificial chemistry follows a sharp phase transition at the "edge of chaos." These claims are supported by simulation-based experiments and are quantitatively validated in the results section (Sec. 5).

### 2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

**Answer:** Yes

**Justification:** A dedicated Limitations section (Sec. 6.2) explicitly discusses modeling assumptions such as the binary monomer system, static environment, simplified fitness proxy, and absence of ecological or selective mechanisms.

### 3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

**Answer:** NA

**Justification:** The paper does not contain formal theoretical results or mathematical proofs; it is based entirely on simulation-based experiments.

### 4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

**Answer:** Yes

510 **Justification:** All key implementation details (e.g., grid size, simulation length, parameter ranges, trial  
511 count) are provided in Sec. 4. Additionally, a reproducibility statement confirms intent to release the  
512 source code and full AI chat logs.

513 **5. Open access to data and code**

514 Question: Does the paper provide open access to the data and code, with sufficient instructions to  
515 faithfully reproduce the main experimental results, as described in supplemental material?

516 **Answer:** Yes

517 **Justification:** The authors commit to releasing source code, configuration files, and chat logs in an  
518 open-access repository upon publication (see Reproducibility Statement).

519 **6. Experimental setting/details**

520 Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters,  
521 how they were chosen, type of optimizer, etc.) necessary to understand the results?

522 **Answer:** Yes

523 **Justification:** The experimental setup is thoroughly described in Sec. 4, including parameter sweeps  
524 (mutation/fragmentation and diffusion/slithering), grid sizes, number of trials, classification criteria,  
525 and simulation duration.

526 **7. Experiment statistical significance**

527 Question: Does the paper report error bars suitably and correctly defined or other appropriate informa-  
528 tion about the statistical significance of the experiments?

529 **Answer:** Yes

530 **Justification:** Each point in the phase space is evaluated over 5 independent trials to account for  
531 stochastic variability, and results are aggregated into probabilities (e.g., of reaching a high-complexity  
532 state). Although error bars are not shown, stochastic robustness is addressed in Sec. 4.4 and Sec. 5.

533 **8. Experiments compute resources**

534 Question: For each experiment, does the paper provide sufficient information on the computer  
535 resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

536 **Answer:** Yes

537 **Justification:** The simulation was run on a standard laptop-class CPU, and experiments are  
538 lightweight—each 50x50 grid simulation runs in under a minute. This is mentioned in the Re-  
539 producibility Statement and will be further detailed in the code repository.

540 **9. Code of ethics**

541 Question: Does the research conducted in the paper conform, in every respect, with the Agents4Science  
542 Code of Ethics (see conference website)?

543 **Answer:** Yes

544 **Justification:** The research complies fully with the Agents4Science Code of Ethics. There are no  
545 human subjects, sensitive data, or dual-use concerns. AI involvement is transparently documented.

546 **10. Broader impacts**

547 Question: Does the paper discuss both potential positive societal impacts and negative societal impacts  
548 of the work performed?

549 **Answer:** Yes

550 **Justification:** The broader impacts are discussed in Sec. 5.3 and Sec. 6, particularly in relation  
551 to astrobiology and the Fermi Paradox. The work contributes to scientific understanding without  
552 presenting foreseeable risks of misuse.