Exploring Flow-Lenia Universes with a Curiosity-driven AI Scientist: Discovering Diverse Ecosystem Dynamics

Thomas Michel*[1], Marko Cvjetko*[2], Gautier Hamon[2], Pierre-Yves Oudeyer[2] and Clément Moulin-Frier[2]

* *Contributed equally*
[1]Univ. Lille, Inria, CNRS, Centrale Lille, CRIStAL, France
[2]Inria Center of the University of Bordeaux, France
firstname.lastname@inria.fr

Abstract

We present a method for the automated discovery of system-level dynamics in Flow-Lenia—a continuous cellular automaton (CA) with mass conservation and parameter localization—using a curiosity-driven AI scientist. This method aims to uncover processes leading to self-organization of evolutionary and ecosystemic dynamics in CAs. We build on previous work which uses diversity search algorithms in Lenia to find self-organized individual patterns, and extend it to large environments that support distinct interacting patterns. We adapt Intrinsically Motivated Goal Exploration Processes (IMGEPs) to drive exploration of diverse Flow-Lenia environments using simulation-wide metrics, such as evolutionary activity, compression-based complexity, and multi-scale entropy. We test our method in two experiments, showcasing its ability to illuminate significantly more diverse dynamics compared to random search. We show qualitative results illustrating how ecosystemic simulations enable self-organization of complex collective behaviors not captured by previous individual pattern search and analysis. We complement automated discovery with an interactive exploration tool, creating an effective human-AI collaborative workflow for scientific investigation. Though demonstrated specifically with Flow-Lenia, this methodology provides a framework potentially applicable to other parameterizable complex systems where understanding emergent collective properties is of interest.

Introduction

Understanding the emergence of complex, system-level properties in self-organizing systems represents a fundamental scientific challenge across many scientific fields (e.g. artificial life, evolutionary biology, and complex systems). From ecological dynamics to open-ended evolution, these phenomena arise from interactions between multiple entities over time, making them difficult to discover through traditional exploration methods that focus on individual patterns or behaviors.

Searching for emergent ecosystem dynamics in artificial systems entail two main challenges.

The first challenge is the simulation of a dynamical system with the potential to self-organize complex interactions and emergent evolution under certain conditions. Computational models offer controlled environments to investigate such emergent properties, with cellular automata (CA) serving as powerful substrates for exploring self-organization and complexity [Langton, 1995, Beer, 2014]. In this paper, we rely on Flow Lenia [Plantec et al., 2023], a recent continuous CA integrating mass conservation and parameter localization. This approach enables the simulation of multiple interacting patterns within a shared environment, creating the potential for emergent evolution and system-wide phenomena such as those shown in Figure 1.

However, the vast parameter spaces of these complex systems brings a second significant methodological challenge: how can researchers efficiently discover diverse emergent dynamics? Random or exhaustive sampling proves ineffective in large parameter spaces, often undersampling low-volume regions that might support interesting behaviors, while manual human exploration cannot systematically cover high-dimensional spaces [Etcheverry et al., 2020, Etcheverry, 2023]. Diversity search algorithms have been proposed as efficient exploration methods in such spaces and have been recently applied to explore large rule spaces of parameterizable cellular automata [Hamon et al., 2024, Faldor and Cully, 2024]. In this paper, we rely on Intrinsically Motivated Goal Exploration Processes (IMGEPs) [Forestier et al., 2022], a family of autotelic diversity search algorithms: IMGEP algorithms self-generate and try to achieve and learn diverse goals to explore a space of behaviors, cf. left side of Figure 2. IMGEPs have been shown to drive efficient exploration of vast spaces of behaviours in multiple domains spanning robotics [Cully et al., 2015, Forestier et al., 2022], biological networks [Etcheverry et al., 2024], code [Wang et al., 2023, Faldor et al., 2024, Pourcel et al., 2024], within the broader family of so-called "AI scientists" systems [Lu et al., 2024, Gottweis et al., 2025]. As they include a general definition of goals as abstract rewards functions subject to arbitrary

This paper features a companion website with a visualization tool that allows to explore a subset of the Flow-Lenia environments explored in the experiments of this paper. The website is available at https://developmentalsystems.org/Exploring-Flow-Lenia-Universes/

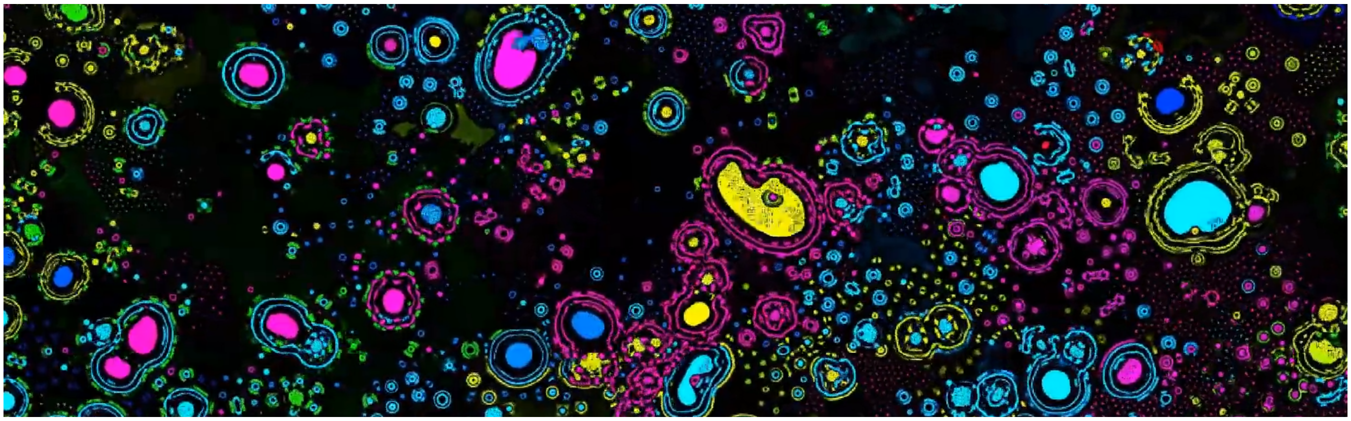

Figure 1: Snapshot of an advanced state of Flow Lenia. Each color represents a different set of localized parameters, determining how matter behaves. Regions with brighter colors represent higher concentrations of mass. In this paper, we present an AI Scientist method, based on IMGEPs, to systematically discover novel dynamics in large-scale Flow Lenia simulations.

constraints, IMGEPs include as a special instance quality-diversity algorithms used in the artificial life community [Faldor and Cully, 2024]. Previous works have shown how IMGEPs enable the discovery of diverse patterns in continuous cellular automata [Reinke et al., 2020, Etcheverry et al., 2021], including emergent sensorimotor agents [Hamon et al., 2024]. However, these approaches focused on discovering individual patterns (spatially localized or Turing patterns). Here we apply it for the first time in Flow Lenia in order to explore a space of simulation-wide metrics that capture aspects of evolutionary dynamics and complexity across multiple scales.

We demonstrate how this approach can systematically discover diverse ecosystem dynamics that would remain hidden using random-search methods. We modified aspects of Flow Lenia to better support multi-species evolution, creating an ideal testbed for our methodology. The resulting discoveries suggest ecological interactions such as feeding, tightly packed clumps of matter resembling colonies, self-organized patterns of different scales in a shared environment and more, as illustrated in Figures 3 and 6.

In addition, we introduce a workflow to bridge automated discovery with human analysis through an interactive exploration tool that provides researchers an effective way to detect promising regions of parameter space. This human-AI collaborative workflow allows scientists to investigate complex emergent phenomena without being limited by either manual parameter tuning or fully automated search, creating a more efficient path for scientific discovery.

Related Works

Our research intersects open-ended evolution, cellular automata, and intrinsically motivated exploration for automated scientific discovery, building upon existing work in these areas.

Open-Ended Evolution in Artificial Systems

Creating artificial systems capable of open-ended evolution remains a fundamental challenge in ALife. [Bedau and Packard, 1996] established quantitative measures of evolutionary activity that have become foundational for evaluating evolutionary dynamics. [Taylor, 2015] and [Soros and Stanley, 2014] identified key requirements for open-ended evolution, including large state spaces, ecological interactions, and hierarchical organization. Despite these advances, most artificial evolutionary systems reach complexity plateaus [Stanley, 2019].

The evaluation of open-endedness has evolved from early parameters like Langton's λ [Langton, 1990] and Wolfram's classifications [Wolfram, 1984] to more comprehensive metrics, such as those in the MODES toolbox [Dolson et al., 2019] and compression-based approaches [Zenil et al., 2015], which measure diversity, complexity, and evolvability in different ways.

Cellular Automata and Flow Lenia

Cellular automata have long served as testbeds for studying open-ended evolution, from Conway's Game of Life [Adamatzky, 2010] to Langton's self-replicating loops [Langton, 1984] and Beer's exploration of autopoiesis [Beer, 2004]. Sayama's work on evolving ecosystems [Sayama, 1999] and genetic evolution in cellular automata [Salzberg and Sayama, 2004] pioneered approaches for encoding and propagating information.

Lenia [Chan, 2018] extended the Game of Life into the continuous domain, generating diverse lifelike patterns. Several variants exist aimed to promote open-ended evolution, such as Chan's model with localized parameters (akin to genes) [Chan, 2023], and Flow Lenia, which, in addition to localized parameters, also incorporates mass conservation, shown to increase evolutionary activity [Hickinbotham and Stepney, 2015].

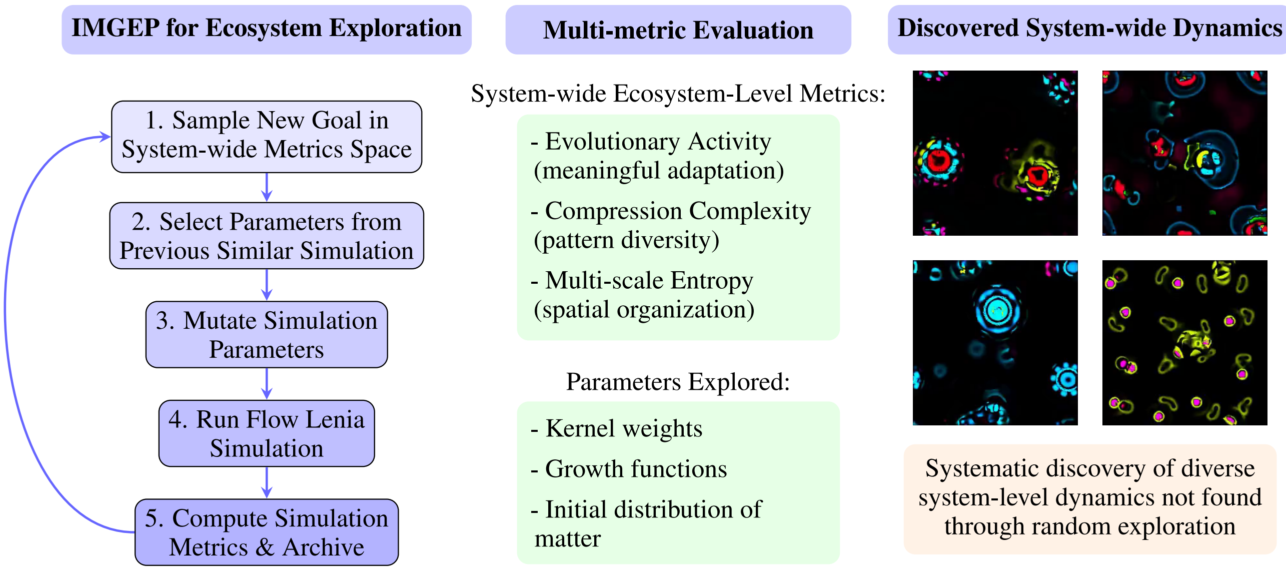

Figure 2: IMGEP approach for exploring a diversity of self-organization in Flow Lenia. The algorithm samples new goals in a system-wide metric space, selects parameters from previous similar simulations, mutates them, runs simulations, and computes metrics to guide further exploration. This process discovers a variety of dynamics, including dense matter clusters akin to colonies, behavior resembling allopatric speciation, feeding, and more.

Intrinsically Motivated Exploration for Automated Scientific Discovery

Random sampling is inadequate to efficiently map parameter space to system's behaviors in complex systems, due to the high dimensionality of the parameter space and the strong non-linearity of the parameter-to-behavior mapping [Etcheverry et al., 2020]. Approaches based on computational models of curiosity or intrinsically motivated exploration [Oudeyer et al., 2007], related to novelty search methods [Lehman and Stanley, 2011], can be used for automated scientific discovery [Etcheverry et al., 2021, 2024], within the large family of AI scientist approaches [Lu et al., 2024, Gottweis et al., 2025]. The IMGEP framework [Forestier et al., 2022] implements a form of intrinsically motivated exploration that efficiently discovers diverse behaviours: it enables to build agents that can self-generate, self-select and pursue their own goals, i.e. autotelic agents ([Colas et al., 2022]). They have demonstrated to be effective approaches for discovering diverse patterns in self-organizing artificial and biological systems [Reinke et al., 2020, Etcheverry et al., 2021, 2024], in physical systems [Falk et al., 2024], or in chemistry [Grizou et al., 2020]. These methods, as well as related methods like novelty-search [Kumar et al., 2024] or quality-diversity [Faldor and Cully, 2024], have proven valuable for exploring self-organization of sensorimotor 'agents' in CA systems such as Lenia [Hamon et al., 2024]. However, they have not previously been applied to investigate evolutionary and/or ecosystemic dynamics in large CA environments.

Measuring Complexity and Evolutionary Activity

Defining and measuring complexity has proven to be a difficult task, with many different metrics proposed over the years [Mitchell and Toroczkai, 2010]. Complexity measures based on information theory [Zenil et al., 2015, Cisneros et al., 2019] connect to Kolmogorov complexity, while concepts like thermodynamic depth [Lloyd and Pagels, 1988] quantify the difference between a system's fine-grained and coarse-grained entropy. [Droop and Hickinbotham, 2012] introduced non-neutral evolutionary activity metrics suited for systems with intrinsic fitness, measuring meaningful adaptive changes while filtering out neutral drift.

We combine complexity and evolutionary activity metrics in order to discover Flow Lenia universes, that not only exibit complex patterns, but also support long-term evolutionary trajectories through local parameter mutations. Our approach aligns well with IMGEP, which helps us systematically illuminate the multi-dimensional metric space, enabling the discovery of new universes and the study of conditions that foster open-ended evolution.

Methods

This section outlines our approach to investigating system-level dynamics in Flow Lenia. We detail the automaton's core components, including its parameter propagation mechanisms, and introduce system-wide metrics designed to identify ecosystem-like environments. The section concludes with a description of our IMGEP implementation.

Flow Lenia

Flow Lenia extends the continuous cellular automaton Lenia by incorporating principles of mass conservation [Plantec et al., 2023]. While Lenia's state space is confined to the unit range $[0,1]$, Flow Lenia expands this to $\mathbb{R}^C_{\geqslant 0}$, where C is the number of channels. This extension allows for unbounded positive real values, reflecting the system's interpretation of cell states as matter density rather than abstract activation levels.

The core components of Flow Lenia include:

- **State space**: $A^t : \mathcal{L} \to \mathbb{R}^C_{\geqslant 0}$ represents the distribution of matter across the grid $\mathcal{L}$ over C different channels
- **Convolution kernels**: $K = \{K_i : \mathcal{L} \to [0,1] \mid i = 1,\ldots,|K|\}$ define the range and strength of interactions between cell states.
- **Growth functions**: $G = \{G_i : [0,1] \to [-1,1] \mid i = 1,\ldots,|K|\}$ determine how interactions influence matter movement.

The system's dynamics unfold in two steps. First, an affinity map U^t is computed for each channel j:

$$U_j^t(x) = \sum_{i=1}^{|K|} h_i \cdot G_i(K_i * A^t_{c_0^i})(x) \cdot [c_1^i = j] \tag{1}$$

where $h_i \in \mathbb{R}$ weights the contribution of each kernel-growth function pair, $K_i * A^t_{c_0^i}$ is the convolution of kernel K_i with its source channel c_0^i, and $[c_1^i = j]$ is the Iverson bracket, equaling 1 when $c_1^i = j$ and 0 otherwise.

Second, a flow field F^t is derived from this affinity map, which determines how matter will move:

$$\begin{cases} F_i^t = (1-\alpha^t)\nabla U_i^t - \alpha^t \nabla A^t_\Sigma \\ \alpha^t(p) = [(A^t_\Sigma(p)/\theta_A)^n]_0^1 \end{cases} \tag{2}$$

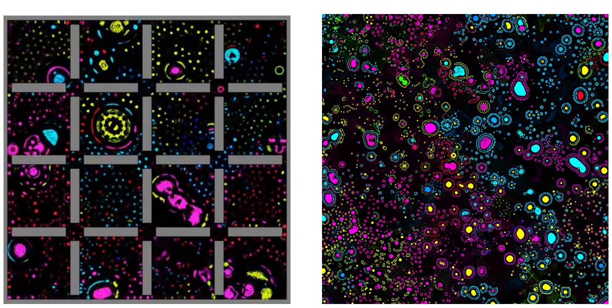

Figure 3: Examples of environments with multiple parameters co-existing and interacting in Flow Lenia. Gray bars represent walls blocking the flow of matter, creating environmental constraints that influence the system's dynamics.

This flow combines an attraction toward high-affinity regions with a diffusion effect (modulated by θ_A) that prevents excessive matter concentration.

To move matter according to the computed flow field while preserving mass conservation, Flow Lenia uses a reintegration tracking method [Moroz, 2020]. This approach works like a grid-based particle system where each cell sends matter according to the flow field. When matter from a source cell flows to a destination that overlaps multiple grid cells, the matter is distributed proportionally across these cells. This ensures that the total amount of matter in the system remains constant, a crucial property that distinguishes Flow Lenia from traditional cellular automata.

Parameter Localization and Mixing Rules

Since Flow Lenia's computation displaces individual cells of matter, it enables one to attach any information to pieces of matter. For instance, it is possible to attach parameters of update rules, i.e. functions which determine the next state of a cell. These parameters will flow with matter, enabling the co-existence of distinct patterns, each driven by different rules, within a shared grid. We define a parameter map $P : \mathcal{L} \to \Theta$, where Θ is the parameter space. In our implementation, we embed the kernel weight vector $h \in \mathbb{R}^{|K|}$, modifying the affinity map computation:

$$U_j^t(x) = \sum_{i=1}^{|K|} P_i^t(x) \cdot G_i(K_i * A^t_{c_0^i})(x) \cdot [c_1^i = j] \tag{3}$$

Because streams of matter carrying different parameters can move to the same destination cell, we need to define a method to resolve conflicts and determine the resulting parameters. We refer to these methods as "mixing rules", and they play a crucial role in the dynamics of the system. The original approach proposed by [Plantec et al., 2023] used stochastic selection, where parameters were chosen randomly weighted by incoming matter quantity. A mixing rule $\mathcal{M}$ is defined as:

$$P^{t+dt}(x_{\text{dest}}) = \mathcal{M}((A^t(x_{\text{src}}), P^t(x_{\text{src}}), I(x_{\text{src}}, x_{\text{dest}})) \mid x_{\text{src}} \in \mathcal{L}) \tag{4}$$

where $I(x_{\text{src}}, x_{\text{dest}})$ denots the proportion of matter flowing from source cell x_{src} to destination cell x_{dest}.

We introduce a "negotiation rule", which considers not only the quantity of incoming matter but also its affinity for the destination environment. The probability of a destination cell adopting parameters from a particular source under the negotiation rule is given by:

$$\mathbb{P}\left[P^{t+dt}(x_{\text{dest}}) = P^t(x_{\text{src}})\right] = \frac{e^{\beta A^t(x_{\text{src}}) I(x_{\text{src}}, x_{\text{dest}}) V^t(x_{\text{src}})}}{\sum_{x \in \mathcal{L}} e^{\beta A^t(x) I(x, x_{\text{dest}}) V^t(x)}} \tag{5}$$

where β is an inverse temperature parameter controlling selection pressure, and $V^t(x_{\text{src}})$ is an affinity map specifically computed for the mixing step:

$$V^t(x_{\text{src}}) = \sum_{j=1}^{C} \sum_{i=1}^{|K|} Q_i^t(x_{\text{src}}) \cdot G_i(K_i * A^t_{c_0^i})(x_{\text{src}}) \cdot [c_1^i = j] \tag{6}$$

Here, $Q_i^t(x)$ is a separate set of parameters used only for the mixing affinity computation.

To promote diversity and evolution, we implement a mutation mechanism at a specified frequency, randomly selecting areas of the grid and applying multivariate Gaussian noise to the parameter map attached to matter.

$$P^{t+dt}(x) = P^t(x) + \varepsilon, \quad \varepsilon \sim \mathcal{N}(0, \Sigma) \tag{7}$$

System-level Metrics

To quantify emergent dynamics in Flow Lenia, we use the following set of metrics:

Non-neutral Evolutionary Activity The non-neutral evolutionary activity metric [Droop and Hickinbotham, 2012] measures population changes in evolutionary systems while filtering out neutral drift:

$$A = \sum_{i} \sum_{t=1}^{T} \Delta_i(t) \tag{8}$$

where i indexes all components (species or genomes) that existed during the simulation, T is the total number of time steps, and:

$$\Delta_i(t) = \begin{cases} (p_i(t) - p_i(t-1))^2 & \text{if } p_i(t) > p_i(t-1) \\ 0 & \text{otherwise} \end{cases} \tag{9}$$

In our case, $p_i(t)$ represents the proportion of simulated mass carrying a local parameter i at time t. We only consider identical local parameters to be the same "genome" type.

Compression-based Complexity Metric We use the size of MP4 video files encoding the simulation as a proxy for complexity, inspired by the concept of Kolmogorov complexity. This approach captures both spatial and temporal patterns, with the MP4 compression algorithm serving as an approximation of the minimal description length.

Multi-scale Matter Distribution To capture the spatial organization at different scales, we compute the entropy of matter distribution at various resolutions. For a given state S, we define a series of downscaled representations $S^1, S^2, ..., S^n$ and compute the entropy for each:

$$H_i = -\sum_{x,y} p_i(x, y) \log p_i(x, y) \tag{10}$$

where $p_i(x, y)$ is the proportion of matter at position (x, y) in the downscaled representation S^i. This provides insights into hierarchical organization, from fine-grained local patterns to large-scale global structures.

IMGEP Implementation

We employ an Intrinsically Motivated Goal Exploration Process (IMGEP) algorithm, as outlined in Algorithm 1. The IMGEP implementation used here iteratively samples target behavioural descriptors as goals, then searching for Flow Lenia parameters that lead to these goals (using mutations over known parameters that are closest to these goals). This progressively leads to discovering and illuminating Flow Lenia simulations that are diverse in that space of descriptors (here system-level metrics). Our parameter space includes parameters related to kernels, growth functions and local rules, along with environmental factors such as the strength of diffusion and mutation rate.

Algorithm 1 IMGEP for Exploring System-wide Properties in Flow Lenia

```
Require: Length of the bootstrapping phase N
Initialize an empty database of discoveries
for each iteration do
    if Number of trials < N then
        Sample random initial parameters
    else
        Sample random goal in system-level metrics space
        Find closest explored goal and its parameter
        Mutate the parameters
    end if
    Set up Flow Lenia with the chosen parameters
    Run Flow Lenia simulation
    Compute achieved goal from simulation results
    Add new parameters and achieved goal to database
end for
```

To address the computational intensity of long-term simulations, we utilized parallel simulations with reduced grid sizes, exploring a broader range of configurations within computational constraints.

Experiments

Our experiments aim to uncover a wide range of behaviors, as measured by the metrics that define the IMGEP goal space. We conduct two experiments: the first investigates Flow Lenia's ability to exhibit ecosystem-like dynamics, while the second focuses on discovering movement patterns of small amounts of matter. Together, these experiments demonstrate the versatility of our method. We also conduct a qualitative analysis of both experiments to highlight the rich dynamics that Flow Lenia can produce. To facilitate this, we develop an interactive exploration tool[1], which enables examination of the relationship between parameter configurations, reached goals, and videos generated from intermediate simulation steps.

Both experiments are conducted on a 256x256 grid, with simulation running for 10,000 time steps.

Exploring Ecosystem-like dynamics

The first experiment uses IMGEP to discover Flow Lenia universes through the metrics described in the previous section. Figure 4 shows reached goals of our method compared to random search. The difference is most evident in evolutionary activity and multi-scale entropy metrics.

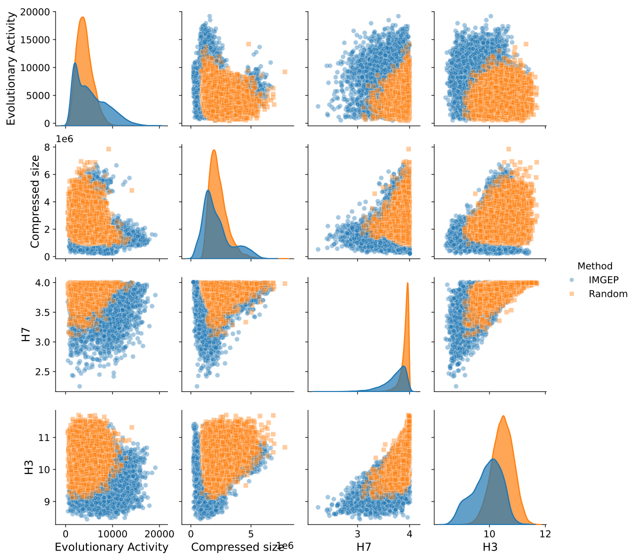

Figure 4: Goal space coverage comparison between IMGEP and random exploration, showing that IMGEP consistently illuminates more of the goal space than random search.

To evaluate the effectiveness our method, we compute the average pairwise distance in the goal space, finding that IMGEP achieves a distance of 0.874, significantly higher than the 0.476 obtained through random search. Additionally, we compute a binning-based coverage by discretizing each goal-space dimension into five bins and counting the number of non-empty bins. IMGEP identifies 576 non-empty bins, in contrast to only 205 discovered via random sampling.

The changes in exploration metrics over time (Figure 5) show that IMGEP found a subset of parameters that strongly affect evolutionary activity. Furthermore, the ongoing increase in coverage suggests many areas of the metric space have yet to be explored.

A closer examination of the videos in the discovery archive (sample frames shown in Figure 6), reveals a striking diversity of self-organizing behaviors, both spatially and temporally. Some universes show the spread of a dominant local parameter that gradually takes over all matter, leading to convergence to an orderly state. Others exibit chaotic dynamics, where matter continuously merges, separates and mutates, with older patterns fading as new ones emerge. In some cases, large, cohesive structures form from multiple localized parameters types, while others resemble dense clusters akin to bacterial colonies. Notably, we also observe behaviors reminiscent of feeding, where larger patterns composed of several localized parameters consume smaller ones (see Figure 7).

Metric	IMGEP	Random Search
Avg Pair. Dist.	**8.74e-01**	4.76e-01
Coverage	**576**	205
EA	**3.44e+03**	1.70e+03
MP4	**1.20e+06**	7.60e+05
H7	**2.26e-01**	8.62e-02
H6	**4.10e-01**	1.66e-01
H5	**5.43e-01**	2.54e-01
H4	**5.67e-01**	3.44e-01
H3	**5.19e-01**	3.79e-01

Table 1: Exploration performance comparison. Higher values indicate broader exploration of evolutionary dynamics.

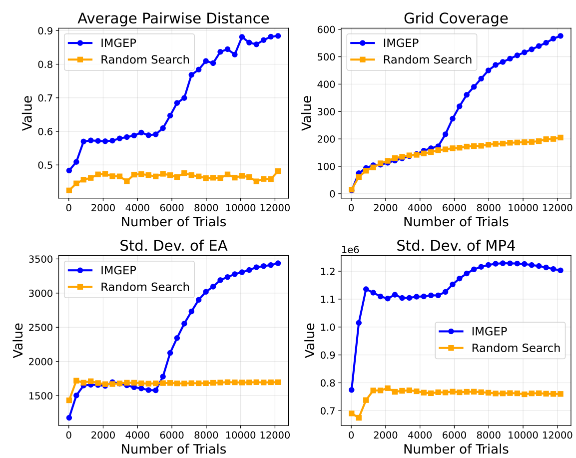

Figure 5: Evolution of exploration metrics over time. The consistent divergence between IMGEP (blue) and random search (orange) indicates IMGEP's sustained ability to discover new evolutionary regimes.

Exploring Matter Movement

The second set of experiments focuses on discovering diverse patterns of matter movement within environments that include obstacles. We design an environment composed of a grid of walls with passages, and place all the initial matter in one corner of the grid (Figure 9, left). We explore such environments with IMGEP through a simple goal space metric: the center of mass of all matter at the end of the simulation

[1]The interactive exploration tool is available at the project website https://developmentalsystems.org/Exploring-Flow-Lenia-Universes/

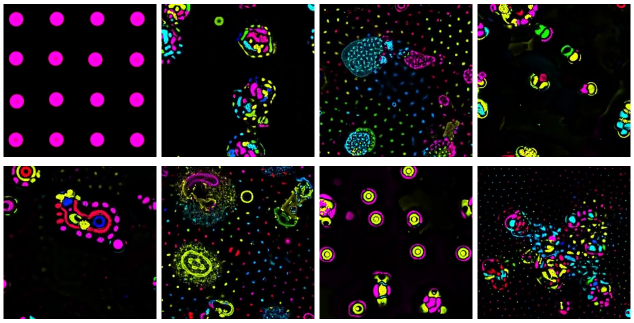

Figure 6: A showcase of discovered diversity while searching for ecosystemic dynamics.

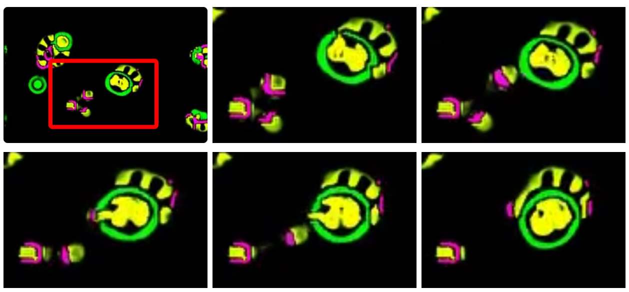

Figure 7: An example of dynamics resembling feeding. The larger, more complex pattern increases it's volume by consuming smaller chunks of matter.

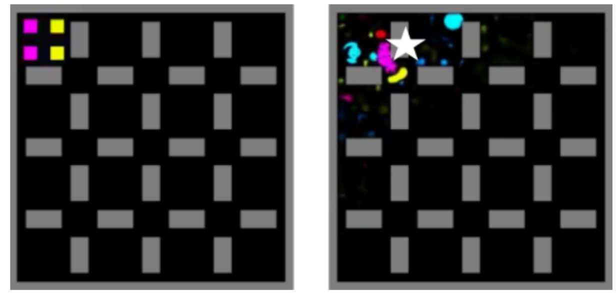

Figure 8: An example of the matter distribution experimental setup. Initial state is shown on the left, and the final state is show on the right, with a white star denoting the center of mass

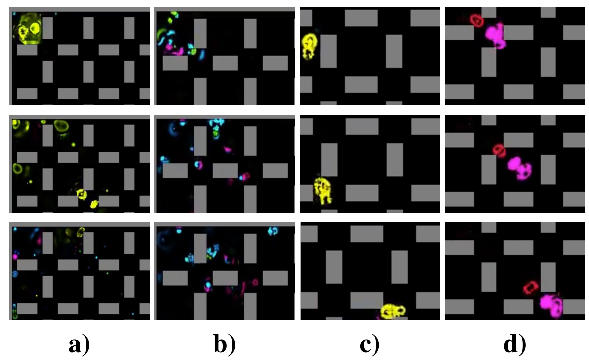

Figure 9: Examples of diverse ways in which matter can move. Each column shows a temporal sequence of: a) Allopatric evolution, b) Initial pattern splitting into agile movers c) pattern turning in a corridor d) direct movement.

(Figure 9, right).

We run IMGEP for 2000 iterations and compare its performance to random exploration. As shown in Figure 10, IMGEP achieves broader coverage in the goal space and reaches points corresponding to matter that has traveled farther from its initial location. This indicates that IMGEP is capable of identifying simulation parameters that enable novel movement behaviors.

The qualitative analysis reveals a rich variety of movement dynamics. We observe both directed and apparently stochastic motion, as well as large patterns fragmenting into smaller, faster-moving entities capable of navigating through narrow corridors. Other cases show dense clusters of matter that gradually diffuse across the environment, with subsequent mutations giving rise to multiple localized rules from a single original pattern-reminiscent of allopatric speciation (see Figure 9). These findings, along with those from the previous experiment, are featured on the project website.

Discussion and Conclusion

We introduced an adaptation of IMGEP diversity search algorithms to explore a broad range of system-level properties, offering a methodical approach for uncovering emergent collective dynamics in complex systems.

While previous applications of IMGEP have focused on discovering diversity of individual CA patterns [Hamon et al., 2024] [Etcheverry et al., 2020], our implementation makes a critical shift to studying multiple interacting patterns potentially governed by different update rules. Rather than optimizing for interesting static properties, we define goals in terms of simulation-wide metrics related to evolutionary activity and complexity. This higher-level application of IMGEP allows us to explore conditions that promote interesting long-term dynamics, rather than merely snapshots. We believe our method is not limited to Flow Lenia, rather, it can be applied to other cellular automata, agent-based models, ecological simulations and complex systems. The only requirement is to use metrics which can capture system-level dynamics of the system.

The findings from our experiments suggest that goal-directed exploration can effectively navigate parameter spaces to locate configurations that support diverse system-level properties. This conclusion is supported by coverage metrics, which show that IMGEP illuminates significantly more goal space than random search, and demonstrate potential for continued exploration.

The patterns observed in our simulations demonstrate the high expressiveness of Flow Lenia. Various organizational structures emerge from the automaton's local rules, with

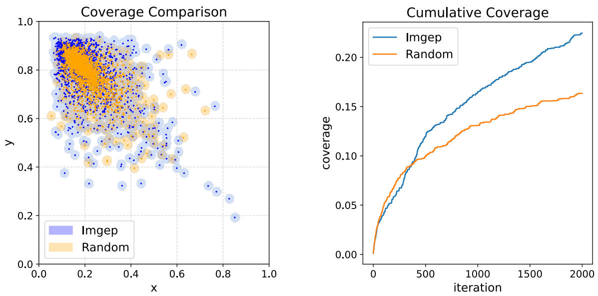

Figure 10: Comparison of reached goals between (left) and cumulative coverage curves (right) between IMGEP and random in the matter distribution experiment.

strong indications of emergent evolutionary and ecological dynamics. We developed an interactive exploration tool which proved to be highly valuable during analyses, allowing us to quickly explore large datasets and uncover diverse dynamics. This tool also helped us assess whether the proposed metrics are correlated with these dynamics. We believe this human-AI collaborative approach fosters more efficient pathways for scientific discovery.

Our multi-metric evaluation framework suggested that the most interesting configurations were not those maximizing any single metric, but rather those exhibiting balance across multiple dimensions. Configurations showing sustained evolutionary activity often featured organization across different spatial scales, maintained parameter diversity despite matter exchange, and demonstrated dynamic yet persistent interactions. This observation aligns with theories that complex behaviors often emerge at the boundary between ordered and chaotic regimes. While these findings remain qualitative, they point to an important direction for future work: developing methods to more rigorously characterize such phenomena.

The limitations of our current implementation include computational constraints that restricted simulation durations, sensitivity to initial conditions that complicated evolutionary analysis, and challenges in definitively distinguishing meaningful adaptations from random variation.

Our investigation of the gathered experimental data has only scratched the surface. In addition to releasing our interactive tool, we will open source the full codebase to facilitate reproducibility and encourage further research. We invite the community to use the released resources to design new metrics, craft richer environments, and develop novel analytical methods for studying ecosystem dynamics in CAs and beyond. Of particular interest is to investigate the rich ways in which emerging patterns interact with each other, and to consider the system from an ecological perspective. Future work could also involve extending simulation times. Applying our method and metrics to other computational models could help identify common conditions that foster complex phenomena across different substrates.

References

Andrew Adamatzky. *Game of life cellular automata*, volume 1. Springer, 2010.

Mark A. Bedau and Norman H. Packard. Measurement of Evolutionary Activity, Teleology, and Life, 1996.

Randall D. Beer. Autopoiesis and Cognition in the Game of Life. *Artificial Life*, 10(3):309–326, July 2004. ISSN 1064-5462. doi: 10.1162/1064546041255539.

Randall D. Beer. The cognitive domain of a glider in the game of life. *Artificial Life*, 20(2):183–206, 04 2014. ISSN 1064-5462. doi: 10.1162/ARTL_a_00125.

Bert Wang-Chak Chan. Lenia-biology of artificial life. *arXiv preprint arXiv:1812.05433*, 2018.

Bert Wang-Chak Chan. Towards Large-Scale Simulations of Open-Ended Evolution in Continuous Cellular Automata. *arXiv preprint arXiv:2304.05639*, 2023.

Hugo Cisneros, Josef Sivic, and Tomas Mikolov. Evolving structures in complex systems. In *2019 IEEE Symposium Series on Computational Intelligence (SSCI)*, pages 230–237. IEEE, 2019.

Cédric Colas, Tristan Karch, Olivier Sigaud, and Pierre-Yves Oudeyer. Autotelic agents with intrinsically motivated goal-conditioned reinforcement learning: a short survey. *Journal of Artificial Intelligence Research*, 74: 1159–1199, 2022.

Antoine Cully, Jeff Clune, Danesh Tarapore, and Jean-Baptiste Mouret. Robots that can adapt like animals. *Nature*, 521(7553):503–507, 2015.

Emily L Dolson, Anya E Vostinar, Michael J Wiser, and Charles Ofria. The modes toolbox: Measurements of open-ended dynamics in evolving systems. *Artificial life*, 25(1):50–73, 2019.

Alastair Droop and Simon Hickinbotham. A quantitative measure of non-neutral evolutionary activity for systems that exhibit intrinsic fitness. In *ALIFE 2012: The Thirteenth International Conference on the Synthesis and Simulation of Living Systems*, pages 45–52. MIT Press, July 2012. doi: 10.1162/978-0-262-31050-5-ch007.

Mayalen Etcheverry. *Curiosity-driven AI for Science: Automated Discovery of Self-Organized Structures*. PhD thesis, Université de Bordeaux, 2023.

Mayalen Etcheverry, Clément Moulin-Frier, and Pierre-Yves Oudeyer. Hierarchically organized latent modules for exploratory search in morphogenetic systems. *Advances in Neural Information Processing Systems*, 33:4846–4859, 2020.

Mayalen Etcheverry, Clement Moulin-Frier, and Pierre-Yves Oudeyer. Hierarchically Organized Latent Modules for Exploratory Search in Morphogenetic Systems, September 2021.

Mayalen Etcheverry, Clément Moulin-Frier, Pierre-Yves Oudeyer, and Michael Levin. Ai-driven automated discovery tools reveal diverse behavioral competencies of biological networks. *eLife*, 13, 2024.

Maxence Faldor and Antoine Cully. Toward artificial open-ended evolution within lenia using quality-diversity. In *ALIFE 2024: Proceedings of the 2024 Artificial Life Conference*. MIT Press, 2024.

Maxence Faldor, Jenny Zhang, Antoine Cully, and Jeff Clune. Omni-epic: Open-endedness via models of human notions of interestingness with environments programmed in code. *arXiv preprint arXiv:2405.15568*, 2024.

Martin J Falk, Finnegan D Roach, William Gilpin, and Arvind Murugan. Curiosity-driven search for novel nonequilibrium behaviors. *Physical Review Research*, 6 (3):033052, 2024.

Sébastien Forestier, Rémy Portelas, Yoan Mollard, and Pierre-Yves Oudeyer. Intrinsically motivated goal exploration processes with automatic curriculum learning. *Journal of Machine Learning Research*, 23(152):1–41, 2022.

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.

Jonathan Grizou, Laurie J Points, Abhishek Sharma, and Leroy Cronin. A curious formulation robot enables the discovery of a novel protocell behavior. *Science advances*, 6(5):eaay4237, 2020.

Gautier Hamon, Mayalen Etcheverry, Bert Wang-Chak Chan, Clément Moulin-Frier, and Pierre-Yves Oudeyer. Discovering sensorimotor agency in cellular automata using diversity search. *arXiv preprint arXiv:2402.10236*, 2024.

Simon John Hickinbotham and Susan Stepney. Conservation of matter increases evolutionary activity. In *European Conference on Artificial Life 2015*, pages 98–105. MIT Press, 2015.

Akarsh Kumar, Chris Lu, Louis Kirsch, Yujin Tang, Kenneth O Stanley, Phillip Isola, and David Ha. Automating the search for artificial life with foundation models. *arXiv preprint arXiv:2412.17799*, 2024.

Chris G Langton. Computation at the edge of chaos: Phase transitions and emergent computation. *Physica D: nonlinear phenomena*, 42(1-3):12–37, 1990.

Christopher G Langton. Self-reproduction in cellular automata. *Physica D: Nonlinear Phenomena*, 10(1-2): 135–144, 1984.

Christopher G. Langton. *Artificial Life: An Overview*. The MIT Press, 07 1995. ISBN 9780262277921. doi: 10.7551/mitpress/1427.001.0001.

Joel Lehman and Kenneth O. Stanley. Novelty Search and the Problem with Objectives. In Rick Riolo, Ekaterina Vladislavleva, and Jason H. Moore, editors, *Genetic Programming Theory and Practice IX*, Genetic and Evolutionary Computation, pages 37–56. Springer, New York, NY, 2011. ISBN 978-1-4614-1770-5. doi: 10.1007/978-1-4614-1770-5_3.

Seth Lloyd and Heinz Pagels. Complexity as thermodynamic depth. *Annals of physics*, 188(1):186–213, 1988.

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.

Melanie Mitchell and Zoltan Toroczkai. Complexity: A guided tour. *Physics Today*, 63:47–, 02 2010. doi: 10.1063/1.3326990.

Mykhailo Moroz. Reintegration tracking. https://michaelmoroz.github.io/Reintegration-Tracking/, 2020. Accessed: 2025-05-05.

Pierre-Yves Oudeyer, Frdric Kaplan, and Verena V Hafner. Intrinsic motivation systems for autonomous mental development. *IEEE transactions on evolutionary computation*, 11(2):265–286, 2007.

Erwan Plantec, Gautier Hamon, Mayalen Etcheverry, Pierre-Yves Oudeyer, Clément Moulin-Frier, and Bert Wang-Chak Chan. Flow-lenia: Towards open-ended evolution in cellular automata through mass conservation and parameter localization. In *Artificial Life Conference Proceedings 35*, volume 2023, page 131. MIT Press, 2023.

Julien Pourcel, Cédric Colas, Gaia Molinaro, Pierre-Yves Oudeyer, and Laetitia Teodorescu. Aces: Generating a diversity of challenging programming puzzles with autotelic generative models. *Advances in Neural Information Processing Systems*, 37:67627–67662, 2024.

Chris Reinke, Mayalen Etcheverry, and Pierre-Yves Oudeyer. Intrinsically Motivated Discovery of Diverse Patterns in Self-Organizing Systems, February 2020.

Chris Salzberg and Hiroki Sayama. Complex genetic evolution of artificial self-replicators in cellular automata. *Complexity*, 10(2):33–39, 2004. ISSN 1099-0526. doi: 10.1002/cplx.20060.

Hiroki Sayama. Toward the realization of an evolving ecosystem on cellular automata. *Proc. Fourth Int. Symp. Artificial Life and Robotics*, pages 254–257, 1999.

Lisa Soros and Kenneth Stanley. Identifying necessary conditions for open-ended evolution through the artificial life world of chromaria. In *Artificial Life Conference Proceedings*, pages 793–800. MIT Press, 2014.

Kenneth O Stanley. Why open-endedness matters. *Artificial life*, 25(3):232–235, 2019.

Tim Taylor. Requirements for open-ended evolution in natural and artificial systems. *arXiv preprint arXiv:1507.07403*, 2015.

Guanzhi Wang, Yuqi Xie, Yunfan Jiang, Ajay Mandlekar, Chaowei Xiao, Yuke Zhu, Linxi Fan, and Anima Anandkumar. Voyager: An open-ended embodied agent with large language models. *arXiv preprint arXiv:2305.16291*, 2023.

Stephen Wolfram. Universality and complexity in cellular automata. *Physica D: Nonlinear Phenomena*, 10(1-2): 1–35, 1984.

Hector Zenil, Fernando Soler-Toscano, Jean-Paul Delahaye, and Nicolas Gauvrit. Two-dimensional kolmogorov complexity and an empirical validation of the coding theorem method by compressibility. *PeerJ Computer Science*, 1:e23, 2015.