Toward Open-Ended Intrinsic Evolution in Lenia Through Multi-Objective Optimization

Niko Lorantos
University of Massachusetts Amherst
nlorantos@umass.edu

Lee Spector
Amherst College
lspector@amherst.edu

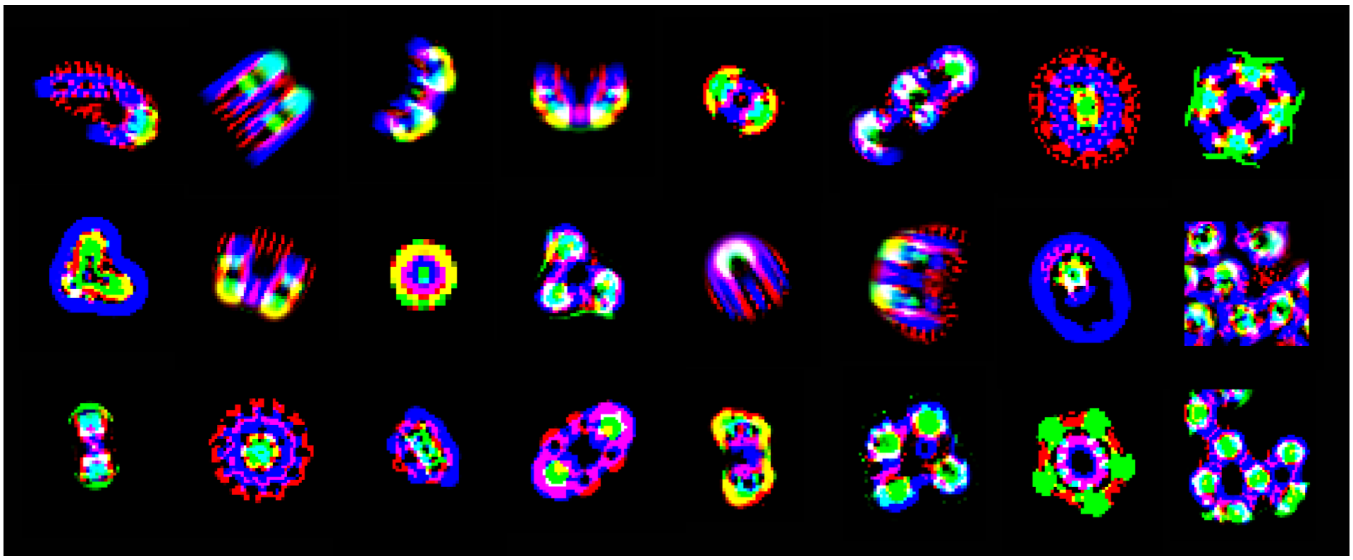

Figure 1: Lenia patterns evolved through intrinsic multi-objective optimization.

ABSTRACT

Artificial life aims to understand the fundamental principles of biological life by creating computational models that exhibit life-like properties. Although artificial life systems show promise for simulating biological evolution, achieving open-endedness remains a central challenge. This work investigates mechanisms to promote unbounded innovation within Lenia, a family of continuous cellular automata, by evaluating individuals against each other with respect to distinctiveness, homeostatic regulation, and population sparsity. These intrinsic fitness objectives encourage the perpetual selection of novel individuals in sparse regions of the descriptor space without restricting the scope of emergent behaviors. We present experiments demonstrating the effectiveness of our multi-objective approach and show that intrinsic fitness objectives allow diverse expressions of artificial life to emerge. We argue that these results are indicative of improved evolutionary dynamics and serve as an important step toward achieving open-ended evolution in artificial systems.

CCS CONCEPTS

• **Computing methodologies → Artificial life**; • **Theory of computation → Bio-inspired optimization**.

KEYWORDS

artificial life, cellular automata, multi-objective optimization, open-endedness, intrinsic evolution

ACM Reference Format:
Niko Lorantos and Lee Spector. 2025. Toward Open-Ended Intrinsic Evolution in Lenia Through Multi-Objective Optimization. In *Proceedings of Genetic and Evolutionary Computation Conference (GECCO'25).* ACM, New York, NY, USA, 4 pages. https://doi.org/10.1145/nnnnnnn.nnnnnnn

1 INTRODUCTION

The emergence of complex lifeforms in nature can be attributed to billions of years of evolutionary innovation. This process relies on the ability to continually produce novelty without limit, a characteristic known as open-endedness (OE). However, achieving OE in artificial life (ALife) systems remains a central challenge that is fundamental to replicating biological complexity and discovering new forms of intelligence. Most artificial systems fail to enable the degree of unbounded innovation and complexification seen in biology, while even the definition of OE remains debated. More concrete

GECCO'25, July 2025, Málaga, Spain

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM
https://doi.org/10.1145/nnnnnnn.nnnnnnn

definitions and robust metrics would be needed to appropriately replicate and classify open-ended behavior in ALife [1, 11, 12].

This work investigates how multi-objective optimization of intrinsic fitness objectives may serve as a step toward achieving OE in ALife simulations. We rank individuals against each other rather than optimizing for a predefined goal to encourage innovation without restricting the scope of emergent behaviors. We argue that this intrinsic multi-objective approach may enable sustained innovation and open-ended evolution in ALife.

1.1 Lenia and Leniabreeder

Lenia [3], a family of continuous cellular automata, is a promising platform for exploring OE in ALife. Lenia generalizes Conway's Game of Life by using continuous states and differentiable update rules based on convolution. It models complex, life-like patterns and emergent behaviors through iterative applications of these rules on a grid. Lenia patterns have demonstrated capabilities such as self-organization, homeostatic regulation, locomotion, entropy reduction, growth, adaptability, and evolvability [4]. This makes Lenia an ideal testbed for studying OE in ALife.

The Leniabreeder framework [6] further enhances Lenia by integrating Quality-Diversity (QD) algorithms to facilitate the selection of diverse and high-performing patterns. Leniabreeder utilizes a Variational Autoencoder (VAE) to represent patterns in a latent descriptor space, enabling the comparison of phenotypes. The AURORA QD algorithm [7] extends this approach by capturing diversity in the VAE latent space. Leniabreeder's unsupervised learning techniques and automated feature discovery are fundamental to enabling OE, allowing the system to adapt without human intervention.

1.2 Open-Endedness

OE remains challenging to precisely define and measure, as true open-ended behavior will transcend any predefined measures of OE [12]. However, Banzhaf et al. [1] have proposed a framework in which OE can be defined as three types of novelty:

- Type-0: Variation; novelty within the model.
- Type-1: Innovation; novelty that changes the model.
- Type-2: Emergence; novelty that changes the meta-model.

In this paper, we define OE as the continual production of novelty of these types [1, 11, 13]. Our approach aims to promote Type-0 and Type-1 novelty, utilizing Leniabreeder's VAE latent space to capture behavioral diversity within the population. Type-2 novelty remains more difficult to observe, as novelty that changes the meta-model will effectively move outside any predefined measure of OE. Hence, attempting to quantify OE may be trivial [12], and we will instead focus on exploring underlying mechanisms that promote OE.

1.3 Intrinsic Evolution

Systems that utilize internal mechanisms to guide evolution are considered to be intrinsic, as the fitness landscape is shaped by the system's current state rather than progressing toward a predetermined goal. Intrinsic evolution enables complexity and diversity to emerge, mirroring the open-ended dynamics of natural evolution. This approach is promising in its ability to guide systems toward meaningful innovation, and has been shown to promote OE in foundational ALife literature.

Novelty search [8], for instance, abandons objectives and instead rewards behavioral novelty. This approach demonstrates efficacy in tasks such as maze navigation and biped locomotion. By abandoning explicit fitness objectives, novelty search enables a continual increase in complexity. This concept is extended in Flow Lenia, a mass-conservative extension of Lenia, where Plantec et al. [9] enable multi-species simulations to drive intrinsic evolutionary processes through competition and symbiosis. Furthermore, Reinke et al. [10] explore intrinsically motivated goal exploration processes (IMGEPs) in self-organizing systems like Lenia, emphasizing the automated discovery of diverse patterns. Their IMGEP-OGL algorithm utilizes deep auto-encoders and achieves efficiency comparable to pretrained systems, highlighting the power of unsupervised learning in uncovering novel self-organized structures.

Together, these approaches emphasize the potential for evolution driven by intrinsic processes and autonomous exploration to replicate the open-ended creativity of natural systems. Decoupling evolution from fixed objectives has proven to aid the emergence of life-like behaviors in artificial environments. We utilize this philosophy in the design of our multi-objective fitness mechanism.

2 METHODOLOGY

We evaluate individuals with respect to three objectives: homeostasis, distinctiveness, and population sparsity. These objectives capture notions of behavioral diversity within the population and promote individuals with complex internal structure in underexplored regions of the descriptor space. Their intrinsic nature allows the system to evolve in accordance with its own evolutionary activity.

Homeostasis rewards stable artificial lifeforms, comparable to those in biological life. We calculate latent variance of an individual, measuring stability in the latent space representation across n timesteps

$$f_1 = -\frac{1}{n}\sum_{i=1}^{n} |\vec{z}_i - \bar{\vec{z}}| \tag{1}$$

where $\vec{z}_i$ are latent encodings and $\bar{\vec{z}}$ is their mean. *Distinctiveness* encourages novelty relative to preexisting individuals. We calculate latent mean distance, measuring divergence from average behavior

$$f_2 = |\bar{\vec{z}} - \mathbb{E}[\bar{\vec{z}}]| \tag{2}$$

where $\bar{\vec{z}}$ is the mean latent encoding of an individual. *Sparsity* rewards individuals in underexplored regions of the descriptor space. We calculate the density of the descriptor space using a radial basis function (RBF) kernel

$$f_3 = -\sum_{a \in \text{archive}} \exp\left(-\frac{|\vec{d} - \vec{d}_a|^2}{2\sigma^2}\right) \tag{3}$$

where $\vec{d}$ is the descriptor vector and σ is the kernel width.

We implement a domination count fitness mechanism to rank members of the population against those in the archive — the set of previously generated patterns — based on these objectives. Domination count of an individual x is calculated as

$$d(x) = |\{y \in A \mid y \prec x\}| \tag{4}$$

where A is the set of archived individuals. An individual y is said to dominate x if y is better than or equal to x in all objectives, and strictly better in at least one objective. The final fitness score of an individual is their negative domination count, so solutions dominated by fewer archive members are considered fitter. This approach maintains a diverse set of individuals representing different trade-offs between objectives, thereby fostering greater exploration of the solution space. By ranking individuals against each other, there is no pressure on the system to converge toward a final goal [8].

We compare our multi-objective domination count approach to a single-objective fitness function optimizing for homeostasis. We conducted 50 trials for each approach, evolving 2500 generations with batch sizes of 256 and a repertoire size of 1024. We recorded the average mass, repertoire variance, and complexity of individuals in the final populations. Mass is calculated as the zeroth spatial moment, or total sum of all cell values, representing the total life content in a pattern. Repertoire variance is calculated as the latent variance across solutions for the last timestep, giving a measure of phenotypic diversity. Finally, we calculate complexity as the gzip compression size of phenotypes across all frames. This serves as an approximation of Kolmogorov complexity, the size of the shortest program that can reproduce a pattern's behavior.

3 RESULTS

Our multi-objective approach displayed increased mass and variance with greater compressibility compared to the baseline implementation. Table 1 summarizes the quantitative results from our experiments.

Metric	Baseline	Multi-Objective	Δ
Mass	3.268	**3.292**	+0.73%
Variance	1.093	**1.103**	+0.91%
Complexity	**3.861**	3.820	-1.06%

Table 1: Results of baseline homeostasis and multi-objective fitness mechanisms across all individuals in final repertoire.

These results indicate enhanced evolutionary dynamics, potentially fostering greater OE. The 0.73% increase in average mass suggests that our approach successfully preserved life content and potentially adaptive traits. Similarly, the 0.91% increase in repertoire variance confirms greater diversity within the population. The observed 1.06% reduction in complexity, coupled with increased mass and variance, points to the emergence of patterns exhibiting greater modular internal structure. This reduced compression size underscores that trivial complexification is avoided as the system continually innovates without converging on a fixed goal. Though improvements are modest, the deviation in average mass, variance, and complexity between the two approaches demonstrates statistical significance with $p < 0.001$ using a two-sample t-test. These results highlight the effectiveness of intrinsic multi-objective optimization in enabling greater open-ended behavior.

We observed subtle qualitative differences between the approaches, and attribute diversity in the baseline approach to the AURORA QD algorithm. However, multi-objective optimization yielded patterns with greater structural variety and enhanced internal modularity, displayed in Figure 3. Figure 2 showcases the behavior of such patterns over multiple timesteps. These results suggest that multi-objective optimization of intrinsic objectives may promote greater variation and innovation within the population. By simultaneously optimizing for homeostasis, distinctiveness, and sparsity, the system generates novelty while preserving stability. This balance is essential for sustaining a population of diverse artificial lifeforms.

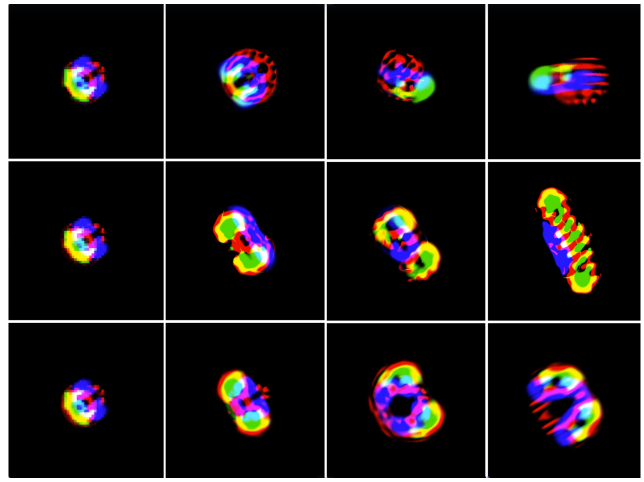

Figure 2: Temporal progression of three patterns starting from the "Aquarium" phenotype (pattern id 5N7KKM) [4, 6].

4 FUTURE WORK

Our findings highlight several promising directions for future research. We intend to apply established measures of open-ended dynamics, such as Bedau and Packard's evolutionary activity test [2] and the MODES toolbox [5], to quantify OE in our system. We posit that more comprehensive measures of OE will offer important insights into the open-ended nature of these systems. Additionally, we aim to explore homeodynamic regulation in place of homeostasis, potentially fostering artificial lifeforms with more intricate and adaptive internal structures. We plan to introduce environmental pressures to assess the system's adaptive complexity, hypothesizing that dominance ranking of intrinsic objectives will outperform a baseline QD algorithm in discovering innovative solutions. Lastly, given that our current experiments capture only the early stages of evolution, we anticipate that extended generational spans could magnify differences in diversity and innovation between the approaches.

5 CONCLUSION

This study demonstrates that intrinsic multi-objective optimization offers a promising step toward achieving open-ended evolution in ALife systems. By designing fitness objectives reminiscent of biological systems, both Type-0 and Type-1 novelty are encouraged. By ranking individuals against each other, we enable greater diversity and exploration of the search space. The theoretical advantages of this approach, coupled with experimental results, indicate that multi-objective optimization can enable innovation by continually driving systems toward novelty.

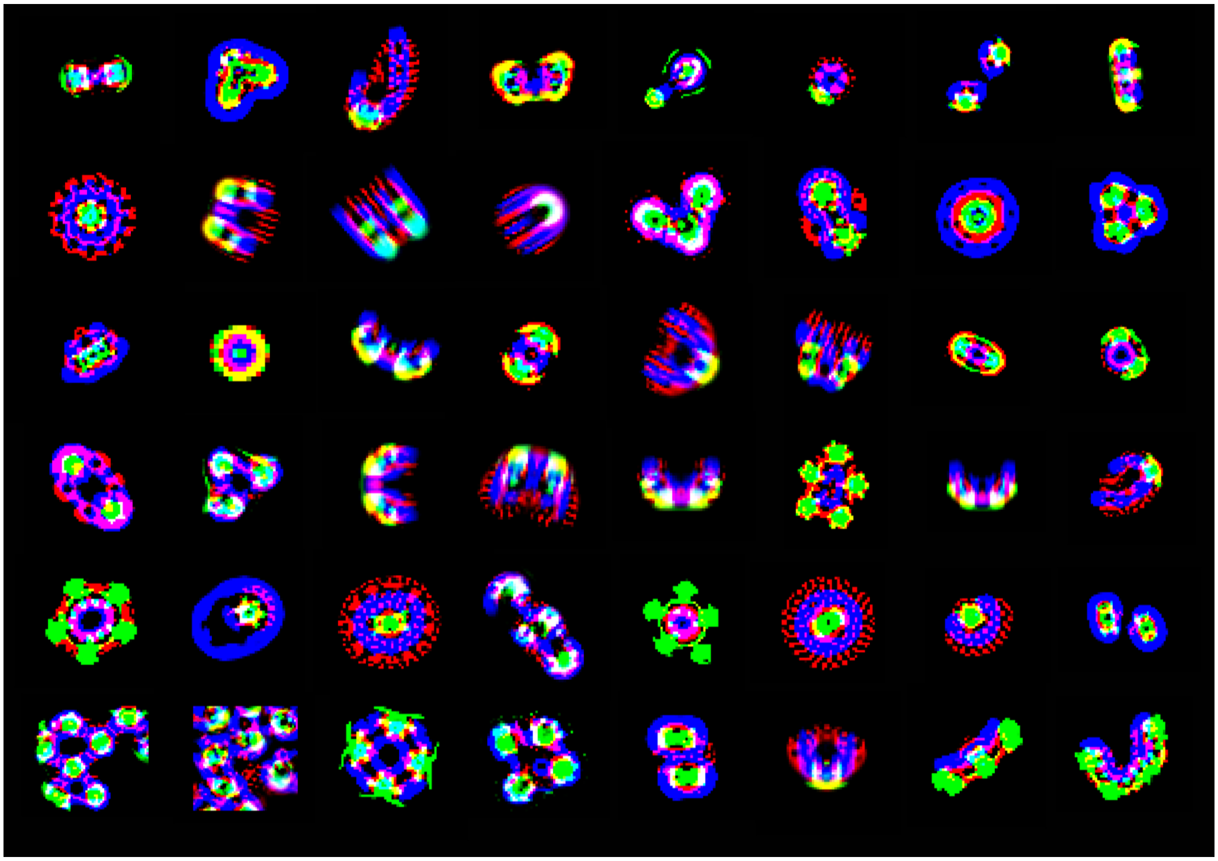

Figure 3: Columns 1-4 display patterns evolved through multi-objective ranking, while columns 5-8 display patterns evolved through homeostasis fitness and the AURORA QD algorithm.

Our contribution extends beyond performance gains to offer a methodological framework for enabling OE by promoting autonomous exploration and decreasing reliance on manual feature design. By aligning selection pressures with intrinsic characteristics of living systems, we foster ongoing innovation akin to natural evolution. Our approach represents a step toward understanding the mechanisms that enable complexity in biological evolution and applying them to artificial systems, bringing us closer to true open-ended behavior in ALife.

REFERENCES

[1] Wolfgang Banzhaf, Bert Baumgaertner, Guillaume Beslon, René Doursat, James A. Foster, Barry McMullin, Vinicius Veloso De Melo, Thomas Miconi, Lee Spector, Susan Stepney, and Roger White. 2016. Defining and simulating open-ended novelty: requirements, guidelines, and challenges. *Theory in Biosciences* 135 (sep 2016), 131–161.

[2] Mark A Bedau, Emile Snyder, C Titus Brown, Norman H Packard, et al. 1997. A comparison of evolutionary activity in artificial evolving systems and in the biosphere. In *Proceedings of the fourth European conference on artificial life*. MIt Press Cambridge, 125–134.

[3] Bert Wang-Chak Chan. 2018. Lenia - Biology of artificial life. arXiv preprint arXiv:1812.05433. arXiv:1812.05433 https://arxiv.org/abs/1812.05433

[4] Bert Wang-Chak Chan. 2020. Lenia and expanded universe. In *Artificial Life Conference Proceedings 32*. MIT Press, Cambridge, MA, USA, 221–229.

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

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

[7] Luca Grillotti and Antoine Cully. 2022. Unsupervised behavior discovery with quality-diversity optimization. *IEEE Transactions on Evolutionary Computation* 26, 6 (mar 2022), 1539–1552.

[8] Joel Lehman and Kenneth O. Stanley. 2011. Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary Computation* 19, 2 (jun 2011), 189–223.

[9] Erwan Plantec, Gautier Hamon, Mayalen Etcheverry, Pierre-Yves Oudeyer, Clément Moulin-Frier, and Bert Wang-Chak Chan. 2023. Flow-Lenia: Towards open-ended evolution in cellular automata through mass conservation and parameter localization. In *Artificial Life Conference Proceedings 35*, Vol. 2023. MIT Press One Rogers Street, Cambridge, MA 02142-1209, USA journals-info …, 131.

[10] Chris Reinke, Mayalen Etcheverry, and Pierre-Yves Oudeyer. 2019. Intrinsically motivated discovery of diverse patterns in self-organizing systems.

[11] Susan Stepney. 2021. Modelling and measuring open-endedness. *Artificial Life* 25, 1 (2021), 9.

[12] Susan Stepney and Simon Hickinbotham. 2024. On the open-endedness of detecting open-endedness. *Artificial Life* 30, 3 (aug 2024), 390–416.

[13] Tim Taylor, Mark Bedau, Alastair Channon, David Ackley, Wolfgang Banzhaf, Guillaume Beslon, Emily Dolson, Tom Froese, Simon Hickinbotham, Takashi Ikegami, and Barry McMullin. 2016. Open-ended evolution: Perspectives from the OEE workshop in York. *Artificial Life* 22, 3 (aug 2016), 408–423.