LLM-Discovered Curvature-Sensitive Growth in Continuous Cellular Automata

Anonymous Authors
Anonymous Affiliation
anonymous@anonymous.org

Abstract

Continuous cellular automata such as Lenia generalize classical reaction–diffusion-like dynamics through the update $A_{t+1} = \mathrm{clip}(A_t + dt \cdot G(K * A_t), 0, 1)$, where K is a radial kernel and G is a pointwise growth function. Existing automated discovery methods search numerically within this fixed functional form. We use a large language model (LLM, Anthropic Claude) as a *generator* of arbitrary symbolic update rules, evaluated under a MAP-Elites quality-diversity loop. The LLM discovers $\cos(4\pi\,\nabla^2 A)$ as a modulator of the growth center—a curvature-sensitive term that is *structurally inexpressible* in the standard $G(K * A)$ family because G receives only the scalar convolution value and is blind to the local Laplacian. The cosine of the Laplacian discretizes a continuous density field into periodic stable bands, producing crystallographic concentric patterns absent from the parametric family even after 125,000 evaluations across 28 physics-derived kernels and 16 growth functions. A controlled three-way ablation (parametric / LLM-guided / random-symbolic) decomposes contributions: the function space supplies behavioral diversity, the parametric template supplies fitness quality, and the LLM supplies precision tuning of structural motifs. The result suggests that AI-guided discovery can extend the mathematical space of self-organizing PDE-like systems beyond hand-designed templates.

1 Introduction

Continuous cellular automata (CCA) sit at the interface of statistical physics and discrete dynamical systems: their continuous-state, continuous-space variants approximate reaction–diffusion equations in a structured Eulerian form, and they exhibit Turing patterns, traveling solitons, and pattern competition familiar from systems such as the FitzHugh–Nagumo and Swift–Hohenberg equations [2]. *Lenia* [1] is a flagship CCA whose update rule

$$A_{t+1} = \mathrm{clip}\big(A_t + dt \cdot G(K * A_t),\ 0,\ 1\big), \tag{1}$$

with radially symmetric kernel K and growth function $G\colon \mathbb{R} \to [-1, 1]$, reproduces over 400 hand-discovered "species" through Gaussian kernels and bell-shaped growth.

Eq. 1 encodes a strong inductive bias: G depends only on the scalar value $u = K * A$ at each pixel and is therefore *blind to the local spatial configuration that produced* u. A flat region and a highly curved region with the same convolution value receive identical growth responses. This separation cleanly mirrors classical reaction–diffusion models (where local reaction depends only on local concentration) but precludes any coupling between growth and spatial derivatives such as $\nabla^2 A$.

Recent automated-search work [3, 4] explores Lenia by perturbing parameters *within* this functional form. We instead use a large language model (LLM) to *generate symbolic update rules* that may freely combine convolutions, Laplacians, and gradients, evaluated by a MAP-Elites quality-diversity

Preprint. Under review.

archive over LZMA pattern complexity and pixel density. Our central finding is that the LLM repeatedly discovers a curvature-sensitive growth modulator, $\cos(4\pi\,\nabla^2 A)$, which is structurally inexpressible in Eq. 1 and produces geometric self-organization with no parametric analogue. We report a three-way ablation that cleanly decomposes *search-space diversity*, *parametric inductive bias*, and *LLM precision tuning* as independent contributions to discovery.

2 Method

Expanded function space (Track A). We expand Lenia's parametric space along three axes: **kernels** from 4 to 28, drawing from molecular and nuclear physics (Yukawa, Lennard-Jones, Morse), fractional calculus (Riesz), signal processing (sinc, Bessel, chirp), and approximation theory (Wendland C^2, sech); **growth functions** from 4 to 16, including reaction–diffusion-derived terms (FitzHugh–Nagumo, Swift–Hohenberg, Brusselator); and **geometric substrates** (flat $\mathbb{R}^2$, sphere S^2, torus T^2, hyperbolic H^2). All kernels are normalised to integrate to 1; all growth functions output $[-1, 1]$.

Open symbolic search (Track B). We use Claude (Anthropic) as a generator of arbitrary update rules combining $\mathrm{conv}(\cdot)$, $\mathrm{laplacian}(\cdot)$, ∇_x, ∇_y, and pointwise functions. Expressions are validated against an operation whitelist, compiled into PyTorch graphs, and evaluated under per-step timeout. The evolutionary loop generates seed expressions, then iteratively recombines parents drawn from the archive (mutation temperature 0.8). The dual-field prompt additionally exposes high-level operations (`curl`, `compete`, `conserve`, `rotate`) that empirically raise hit rates by an order of magnitude over abstract conceptual prompts.

Random-symbolic ablation (Track C). To isolate the LLM's contribution, Track C samples 1,000 symbolic expressions uniformly over the same operation grammar as Track B, with no LLM in the loop.

Behavior space. All three tracks share a MAP-Elites archive indexed by LZMA compression ratio c (Kolmogorov-complexity proxy [6]) and pixel density a, discretised on a 20×20 grid. Only the highest-fitness entry per bin is retained, and entries from all tracks compete for the same bins.

Evaluation protocol. Track A runs at 128^2 resolution for 8,000 steps per evaluation ($\sim$125K total evaluations on $4\times$H100 GPUs); Tracks B and C run at 64^2 for 2,000 steps. Selected entries from all tracks are re-evaluated under *identical* conditions (128^2, 8,000 steps, 10 random seeds) for the controlled comparison reported in Section 3. The fitness is mean LZMA complexity over the trajectory's terminal segment.

3 Results

Archive coverage. The unified archive covers 291/400 bins (72.8%), 285 from Track A and 6 from Track B. Empty bins concentrate in the high-complexity, low-density region where extended textures cannot persist.

The discovered structure. The hero entry from Track B implements a growth function whose center μ is spatially modulated by the local Laplacian:

$$\mu(x, y) = \mu_0 + \varepsilon \cos\big(4\pi\,\nabla^2 A(x, y)\big), \qquad \varepsilon = 0.03,\ \mu_0 = 0.12, \tag{2}$$

combined with a bandpass density window and a secondary harmonic (full expression in supplementary material associated with the project archive[1]). Eq. 2 discretises the continuous Laplacian field into periodic zero crossings, producing concentric crystallographic bands (Fig. 1). A manual ablation confirms that the curvature-sensitive term *alone* (Eq. 2 without the bandpass and harmonic components) produces Turing-like spot patterns (complexity 0.13), while the full expression produces the geometric ring pattern (complexity 0.94). All three components are necessary, indicating that the LLM's contribution includes both the structural innovation of curvature-sensitive growth and the precise arrangement of supporting nonlinearities.

[1] Code and full archive will be released on acceptance.

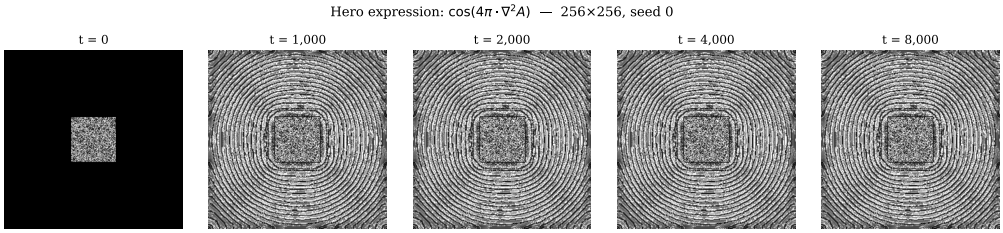

Figure 1: **Emergence of crystallographic bands from a random seed under the discovered curvature-sensitive rule.** A 256×256 noise initial condition (left) self-organises into concentric ring patterns over 8,000 steps. The geometry is set by the resonance between the kernel radius and the 4π frequency in $\cos(4\pi\,\nabla^2 A)$.

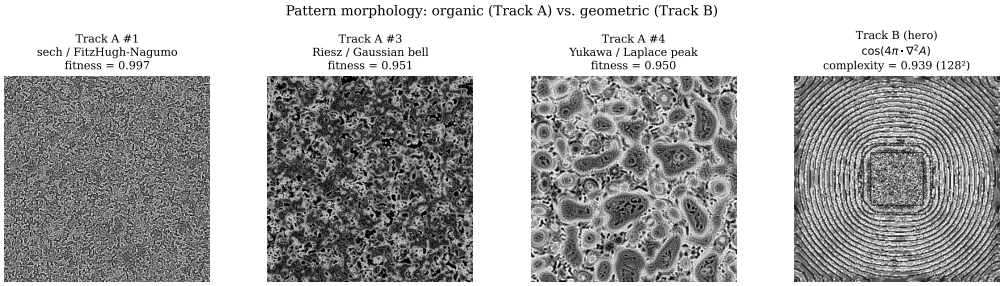

Figure 2: **Morphological divergence between parametric and curvature-sensitive dynamics.** *Left*: top Track A entries produce organic, irregular textures characteristic of reaction–diffusion (sech / FitzHugh–Nagumo, Yukawa / Laplace peak, Riesz / Gaussian bell). *Right*: the LLM-discovered Laplacian-modulated rule produces concentric crystallographic bands. Both shown at 128^2 after 8,000 steps.

Structural inexpressibility. The growth function in Eq. 1 accepts a single argument $u = K * A$. The Laplacian $\nabla^2 A$ is a distinct linear operator that cannot be reconstructed from u alone for any choice of K: $K * A$ encodes a weighted local average, $\nabla^2 A$ a second-derivative trace. Therefore $G(\nabla^2 A,\ K * A)$ lies outside the range of the parametric family regardless of kernel or growth choice. This is not merely a parameter the search did not find—it is a function not present in the function family at all.

Robustness. The discovered rule produces highly consistent patterns across 20 runs at two resolutions: at 128^2, alive fraction 0.977 ± 0.001, complexity 0.939 ± 0.002, 6.1 ± 1.6 clusters; at 256^2, alive fraction 0.983 ± 0.001, complexity 0.899 ± 0.002, 29.3 ± 2.7 clusters. Coefficient of variation across seeds is $< 0.3\%$, comparable to the most reproducible parametric entries.

Three-way ablation. Table 1 compares the three tracks under matched re-evaluation. The expanded parametric family wins fitness rankings, but the LLM track contributes a structurally distinct entry. The random-symbolic baseline (Track C) independently discovers the motif $\cos(\nabla^2 A)$ in 6/136 entries, but *without the* 4π *frequency* that creates spatial resonance with the kernel radius—instead landing on smaller frequencies that yield broader, less crystalline patterns. The LLM's mathematical knowledge enables precision tuning of structural motifs that random sampling discovers in coarser form.

Table 1: Three-track comparison under matched re-evaluation (128^2, 8,000 steps).

Metric	**Track A (parametric)**	**Track B (LLM)**	**Track C (random sym.)**
Evaluations	∼125K	∼5K	1,000
Archive bins	291	6	136
Best fitness	0.997	0.939	0.967
Entries > 0.9	12	1	3

Decomposition. The three tracks suggest three independent contributions to discovery in self-organizing PDE-like systems. *Function space* provides behavioral diversity: 1,000 random symbolic expressions cover 136 of 400 bins, including regions inaccessible to the parametric family. *Parametric inductive bias* provides fitness quality: Track A wins 132 of 136 head-to-head fitness comparisons against Track C, reflecting that the $G(K * A)$ template encodes structural smoothness assumptions that survive long evaluation. *LLM precision tuning* provides qualitative novelty within structurally productive regions: the only structurally distinct entry in the archive (the curvature-sensitive growth) emerges from the LLM track, not from random sampling, and the LLM additionally tunes frequencies and coefficients in a way random sampling does not.

4 Discussion

The result identifies a concrete instance where AI-guided symbolic search extends a self-organizing PDE-like system beyond a hand-designed parametric family along a direction that has clear physical interpretation: coupling growth dynamics to local curvature. Such a coupling is rare but not unknown in physics—it appears, for example, in phase-field models of curvature-driven interface motion—but it is absent from the standard CCA literature, where growth is by construction a function of a smoothed local average.

Connection to physics-style discovery. Our setup parallels recent neural-symbolic and LLM-driven approaches to discovering governing equations [5] but operates in a regime where the "ground truth" is open-ended: there is no fixed dataset to fit, only a behavior space to explore. The closest physics analogue is the discovery of new Turing-instability mechanisms; the analogue we discover differs in that the spatial coupling is not diffusive but *phase-modulating* via a periodic function of the Laplacian.

Limitations. (i) Substrate breadth: we examine only Lenia and have not tested whether the curvature-sensitive motif transfers to Gray–Scott, Cahn–Hilliard, or particle-based systems. (ii) Search budgets are unequal across tracks; the controlled re-evaluation partially compensates but does not match compute. (iii) Track B is warm-started with Track A's archive; without warm start, LLM expressions degenerate at much higher rates. (iv) The LLM (Claude) is a closed-source system, limiting reproducibility of the generation process; the discovered expression itself, however, is fully reproducible from the released code.

Outlook. Two directions are immediately promising. First, the curvature-sensitive motif suggests a small family of operator-valued growth functions $G(K * A, \nabla^2 A, \|\nabla A\|, \dots)$ that can be enumerated and characterised systematically without LLM involvement, mirroring the progression in physics from a single discovered mechanism to a derived family. Second, evaluating the same LLM-driven loop on classical PDE substrates—where "correctness" can be checked against analytical or numerical reference dynamics—would test whether the discovery pattern generalises beyond CCAs to physically grounded systems.

References

[1] Bert Wang-Chak Chan. Lenia – biology of artificial life. *Complex Systems*, 28(3), 2019. arXiv:1812.05433.

[2] M. C. Cross and P. C. Hohenberg. Pattern formation outside of equilibrium. *Reviews of Modern Physics*, 65(3):851–1112, 1993.

[3] Maxence Faldor and Antoine Cully. Toward artificial open-ended evolution within Lenia using quality-diversity. In *Proceedings of ALIFE 2024*. MIT Press, 2024. arXiv:2406.04235.

[4] Akarsh Kumar, Sam Earle, Minqi Jiang, and Kenneth O. Stanley. Automating the search for artificial life with foundation models. *Artificial Life*, 31(3), 2024. arXiv:2412.17799.

[5] Bernardino Romera-Paredes, Mohammadamin Barekatain, Alexander Novikov, Matej Balog, M. Pawan Kumar, Emilien Dupont, Francisco J. R. Ruiz, Jordan S. Ellenberg, Pengming Wang, Omar Fawzi, Pushmeet Kohli, and Alhussein Fawzi. Mathematical discoveries from program search with large language models. *Nature*, 625:468–475, 2023.

[6] Hector Zenil, Santiago Hernandez-Orozco, Narsis A. Kiani, Fernando Soler-Toscano, Antonio Rueda-Toicen, and Jesper Tegnér. A decomposition method for global evaluation of Shannon entropy and local estimations of algorithmic complexity. *Entropy*, 20(8):605, 2018.