Classifying the fractal parameter space of the Lenia Orbium

Ivan Yevenko
iyevenko@uwaterloo.ca

Abstract

The Lenia family of continuous cellular automata can be viewed as the iterative application of a constant function with unpredictable convergence properties. This makes it mathematically analogous to the Mandelbrot set and neural network trainability, which have both been shown to have fractal convergence boundaries. Using an escape-time algorithm, we plot the stability of the *Orbium unicaudatus* species as a function of two parameters at a time, generating fractals that persist on multiple spatial scales. We categorize regions in this parameter space and explore them to find a set of familiar species, one novel specimen, and many non-trivial variations of Orbium that fundamentally rely on discretization to survive. Based on these discoveries, we hypothesize the existence of many complex undiscovered species hidden in the fractal parameter spaces of the rest of the Lenia zoo.

Introduction

Escape-time algorithms were first introduced by Mandelbrot (1982) to visualize the Mandelbrot fractal. Similar approaches have since been used to plot quadratic Julia sets (Julia, 1918), Newton fractals (Barnsley et al., 1988), Burning Ship fractals (Michelitsch and Rössler, 1992) and most recently, the convergence boundary of neural networks (Sohl-Dickstein, 2024). These methods work by iterating a constant function f until $x_n = f^n(x_0)$ satisfies a condition, or the iteration count n reaches an arbitrarily chosen maximum value. An image can then be generated by varying f or x_0 over a 2-D grid (e.g. $x_0 \in \mathbb{C}$ for the Julia set) and coloring each point based on the final value of n.

Motivated by this method of plotting convergence boundaries of iterated functions, we hypothesized that the recurrence relation describing the dynamics of the continuous cellular automaton *Lenia* (Chan, 2019) also has a fractal convergence boundary. *Discrete Lenia* is defined as:

$$\mathbf{A}^{t+1} = \left[\mathbf{A}^t + \frac{1}{T}\, G_{\mu,\sigma}(\mathbf{K}_{R,\beta} * \mathbf{A}^t)\right]_0^1 \qquad (1)$$

Where $\mathbf{K}_{R,\beta}$ is a sum of gaussian bumps weighted by β, with maximum radius R. G is a gaussian parameterized by μ and σ and applied pointwise to the discrete convolution of the kernel $\mathbf{K}_{R,\beta}$ and the current state $\mathbf{A}^t$. Finally, T is the time step or temporal resolution, and $[]_0^1$ is the clipping function $[x]_0^1 = \max(0, \min(1, x))$.

Recently, Davis and Bongard (2022) and Kojima and Ikegami (2023) found that some patterns in Lenia become unstable when T is raised beyond a certain threshold, contradicting the intuition that a more accurate approximation of the underlying continuous differential equation should be more stable. In their subsequent work, Davis (2023) call these patterns *non-platonic* and find patterns that fundamentally rely on discretization across many types of continuous cellular automata.

Methods

In this work, we limit our study to the *Orbium unicaudatus* species, for which we hold $\mathbf{A}^0$ and β at their initial values [1] while varying the T, μ, σ, and R parameters in pairs. To generate the fractals, we linearly sweep over two parameters in a 640×640 grid and run a 64×64 Lenia simulation for each point. The parameter ranges were chosen arbitrarily to showcase particularly interesting regions. Similarly to other escape-time fractals, each pixel's color is determined by the number of steps until the simulation reaches a point of no return. Specifically, we count the number of timesteps t until $\mathbf{A}^t = \vec{0}$. If the simulation runs for $25T$ steps without reaching the zero state, we consider the pattern to be stable.

Observations

Figure 1 shows fractals emerging across all 6 planes of the 4-dimensional parameter space, and across at least 2 decades of scale. We classify the stable regions of the fractals (those colored brightest) in the following three categories:

i) Orbium island: This is a region of (μ, σ, R, T)-space around the original Orbium parameters (shown with a blue star) that contains minor variations of Orbium. The top left of Figure 1 B) resembles parts of the $\mu - \sigma$ map in Figure 9 of Chan (2019), so we believe Orbium island to be a good approximation of Orbium's *niche* (Chan, 2019).

[1] As defined in `https://github.com/Chakazul/Lenia/blob/master/Python/animals.json#L5`

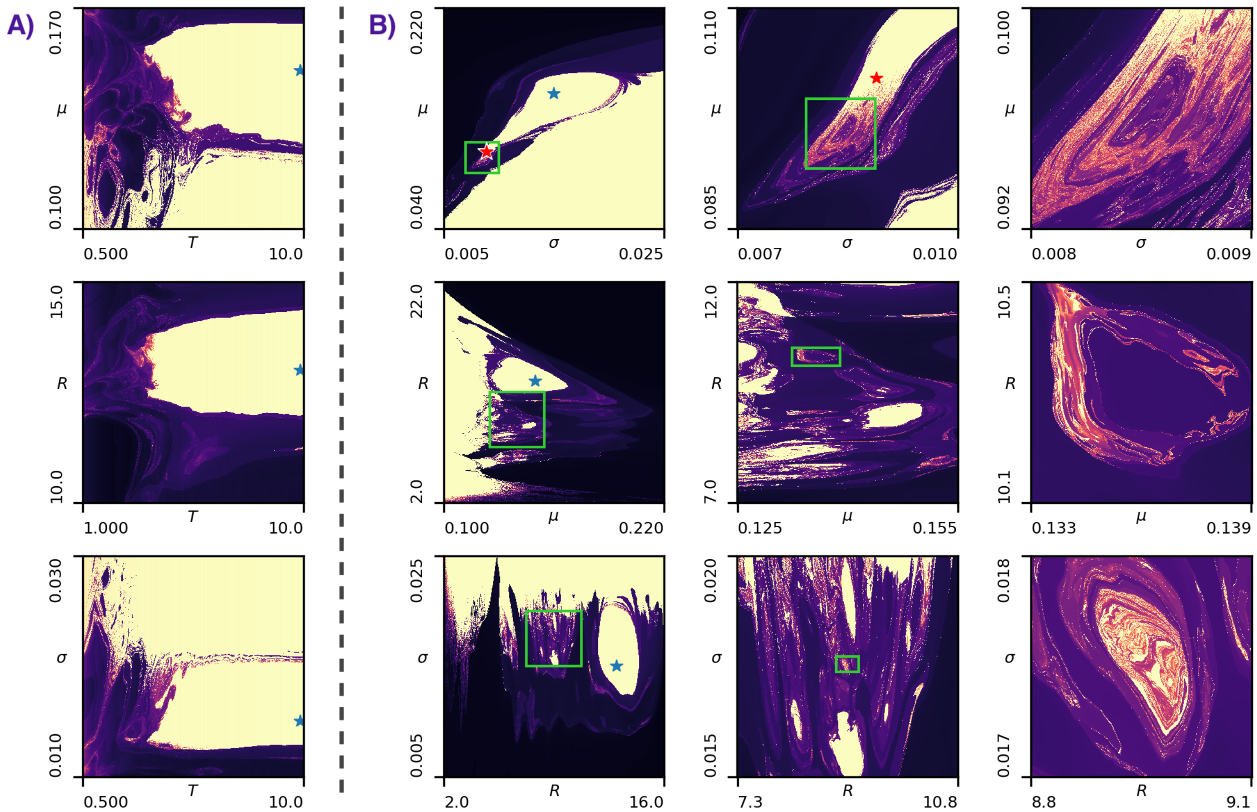

Figure 1: Fractals generated by varying two parameters of *Orbium unicaudatus*. Brightest pixels correspond to a survival time of $25T$ and darkest are ≈ 0. Blue stars mark original Orbium parameters, red stars mark new chaotic glider. **A)** μ-T, R-T and σ-T planes. **B)** μ-σ, R-μ and σ-R planes. Each row contains 3 zooms of the same fractal. Green boxes indicate zoom regions.

ii) Trivial mass: The larger mass adjacent to Orbium island corresponds to parameters that cause the soliton to explode and reach steady state, much like randomly-initialized Lenia. Since the mass of these patterns is nonzero, they are considered stable by our heuristic. This is a subset of the chaotic forest described in (Chan, 2019).

iii) Chaotic boundary: Between the Orbium island and the trivial mass is a chaotic archipelago containing many sporadic patterns. Here, we found non-trivial variations of Orbium that rely heavily on temporal and spatial discretization, as well as the familiar species *Gyrorbium*.

On Orbium Island, we find a continuum of species containing Orbium bicaudatus ($R = 13.3$, $\sigma = 0.0125$), Gyrorbium ($\mu = 0.13436619, \sigma = 0.01792645$), and a new chaotic glider on the long tail of the $\mu - \sigma$ fractal ($\mu = 0.102, \sigma = 0.0089$). This non-platonic glider is most similar to Orbium unicaudatus ignis. In the chaotic boundary, we mostly find wiggling Orbiums ($R = 8.85, \mu = 0.1398$) and chaotic Gyrorbiums ($R = 9.01095, \sigma = 0.02065$).

In Figure 1 A) we observe that under some circumstances, the pattern becomes unstable upon increasing T from a stable point. All else being equal, setting $\sigma = 0.0127$ causes the pattern to be stable for $T = 5$ but unstable for $8 \leq T \leq 100000$, and likely indefinitely. This confirms that non-platonic variants of Orbium exist with coarse enough spatial resolution, and hints at the fact that one might find non-platonic species on the fractal boundary of any otherwise platonic species.

Future Work

The evolutionary methods used to find all existing Lenia species had to traverse a very high-dimensional search space. Our results indicate that imaging low-dimensional slices can also be an effective method for finding new species. Based on this insight, we propose several research directions to build on this work:

1. Can algorithms leveraging knowledge of the parameter space be used to find species more efficiently?
2. What are better heuristics than survival time for identifying complex life forms within the fractals?
3. How do the fractal parameter spaces of all other Lenia species compare? How much do they overlap?

References

Barnsley, M. F., Devaney, R. L., Mandelbrot, B. B., Peitgen, H.-O., Saupe, D., and Voss, R. F. (1988). *The Science of Fractal Images*. Springer New York, NY.

Chan, B. W.-C. (2019). Lenia: Biology of artificial life. *Complex Systems*, 28(3):251–286.

Davis, Q. T. (2023). Discretization-dependent dissolution of gliders in (dis)continuous systems: Non-platonic self-organization in complex systems. *Innovations in Machine Intelligence (IMI)*, 3:1–23.

Davis, Q. T. and Bongard, J. (2022). Step size is a consequential parameter in continuous cellular automata.

Julia, G. (1918). Mémoire sur l'itération des fonctions rationnelles. *Journal de mathématiques pures et appliquées*, 1:47–245.

Kojima, H. and Ikegami, T. (2023). Implementation of lenia as a reaction-diffusion system. In *The 2023 Conference on Artificial Life*, ALIFE 2023. MIT Press.

Mandelbrot, B. B. (1982). *The fractal geometry of nature*, volume 1. WH freeman New York.

Michelitsch, M. and Rössler, O. E. (1992). The "burning ship" and its quasi-julia sets. *Computers & graphics*, 16(4):435–438.

Sohl-Dickstein, J. (2024). The boundary of neural network trainability is fractal.