Emergence of Self-Replicating Hierarchical Structures in a Binary Cellular Automaton

Bo Yang

Adjacent Lab, New York, United States
bo@adjacentlab.com

Abstract

We have discovered a novel transition rule for binary cellular automata (CA) that yields self-replicating structures across two spatial and temporal scales from sparse random initial conditions. Lower-level, shapeshifting clusters frequently follow a transient attractor trajectory, generating new clusters, some of which periodically self-duplicate. When the initial distribution of live cells is sufficiently sparse, these clusters coalesce into larger formations that also self-replicate. These formations may further form the boundaries of an expanding complex on an even larger scale. This rule, dubbed "Outlier," is rotationally symmetric and applies to 2D Moore neighborhoods. It was evolved through Genetic Programming during an extensive search for rules that foster open-ended evolution in CA. While self-replicating structures, both crafted and emergent, have been created in CA with state sets intentionally designed for this purpose, the Outlier may be the first known rule to facilitate nontrivial emergent self-replication across two spatial scales in binary CA.

Traditionally, the "building blocks" one level down from self-replicating structures in cellular automata (CA) are cells with multiple possible states. Notable examples include the 29 states in Von Neumann's original universal construction machine (von Neumann, 1966), and the 8 states in Langton's 86-cell self-replicating loops (Langton, 1984). Both also require carefully designed initial configurations. Emergent self-replicating structures from random initial conditions have previously been achieved using CA with 8-bit state sets (Chou and Reggia, 1997). In all these cases, each state or its subcomponent generally assumes a specific "role," such as signal passage, replication trigger, or structural protection, all of which are vital elements of the intended replication mechanisms.

In contrast, a cell in a binary cellular automaton carries minimal information and is unlikely to perform any specific role. Consequently, if self-replicating structures were to emerge from randomness in binary CA, the "building blocks" must be clusters of cells, emerging solely from the rule itself. In other words, it would be more complex than complicated. This could be of interest to studies of emergence and self-organization. We report the discovery of such a rule, named "Outlier" (figure 1).

The Outlier was serendipitously discovered during an extensive automated search for binary CA rules that would support open-ended evolution (OEE) (Bedau et al., 2000). Genetic Programming (GP) was used in the specific search run that led to Outlier. Each rule used as genotype was represented as a tree structure of bitwise logic operations. "Novelty search" (Lehman and Stanley, 2008) was adopted for fitness measurements. In our implementation, we extract a feature vector, $\mathbf{F}$, for each rule from the complexity profile (Bar-Yam, 2003) of CA bitmaps in the later stages of convergence. For each new rule, a novelty score is calculated based on the distances from $\mathbf{F}$ to its k nearest neighbors in the space of previously computed $\mathbf{F}$, and is used as the fitness. During the novelty search, the Outlier was algorithmically tagged as "novel" enough for visual inspection, where self-replicating patterns were found. However, nothing substantially more complex has been observed thus far.

The Outlier is rotationally symmetric but not mirror symmetric. Similar to many solutions produced by genetic algorithms or GP, it neither has a discernible structure nor a definable formulation. Notably, the "live" bits of its rule table representation are much denser than that of Conway's Game of Life.

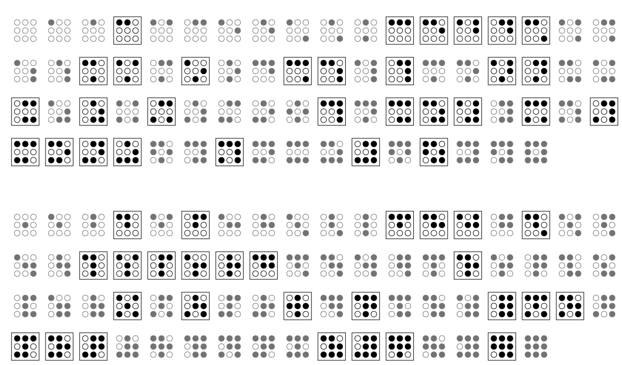

Figure 1: The Outlier Rule. The center cell in each of the boxed neighborhoods and their three quarter-turn rotations stays alive. Filled and empty circles stand for live and dead cells respectively.

Under this rule, from initial random conditions, hierarchical structures often emerge on three spatial scales: clusters of cells, self-replicating formations consisting of multiple clusters, and expanding complexes driven by multiple replicating formations (figure 2). Shapeshifting clusters, each composed of a few dozen live cells at most, form in less than a hundred steps, with most eventually vanishing. In a sparse and sufficiently large grid, a small fraction of the clusters evolve into larger self-replicating formations by spawning new clusters. These formations expand their territory by creating copies of themselves while slowly traversing, until they collide with clusters outside their territory. When a single replicating formation emerges in the middle of a sufficiently large empty area, it periodically generates new formations, forming an still larger structure, or a "complex," at a higher scale. A complex continuously expands until it either occupies all available space or until it collides with other structures outside its territory.

We have also identified temporal loops on two characteristic time scales. These periodic reappearances were observed in experiments that seeded an empty grid with a single, isolated 3 by 3 cluster in the center. Out of all the 140 possible initial configurations, most die out, but two types of seeds evolve into replicating formations by following the same trajectory $\mathcal{A}_s$ thereafter. A detailed examination of $\mathcal{A}_s$ reveals that many clusters periodically reappear, sometimes replicate with a 143-step period, rotating 90 degrees each time. As the formation of clusters expands, entire formations begin to reappear with a period of 1556 steps, some of which are capable of self-replication and are therefore spatially identifiable. Many types of identical clusters are components of these formations, shape-shifting in sync and reappear every 1556 steps.

Under sparse initial random conditions, most arbitrarily formed clusters ultimately disappear. The rare survivors predominantly follow the same trajectory $\mathcal{A}_s$. However, if more than one cluster survives, their individual evolutions along $\mathcal{A}_s$ are derailed upon collision between clusters originating from different seeds. Consequently, the density of the initial random configuration determines the likelihood of replicating formations' existence. When the initial grid is dense, surviving clusters bypass the formations and evolve directly into a semi-chaotic phase, where they continuously change shape, split, merge, and interact with each other.

In summary, a cellular automaton operating under the Outlier rule transitions into one of three phases: empty, semi-chaotic, or replication at the formation level. The latter phase is characterized by a trajectory that resembles an expanding transient attractor. Reappearances of both clusters and formations attach sub-loops to the trajectory, each with distinct characteristic period lengths: 143 steps for clusters, and 1556 steps for formations. However, this attractor is transient as the complex eventually either exhausts the available space for expansion, or collides with other structures,

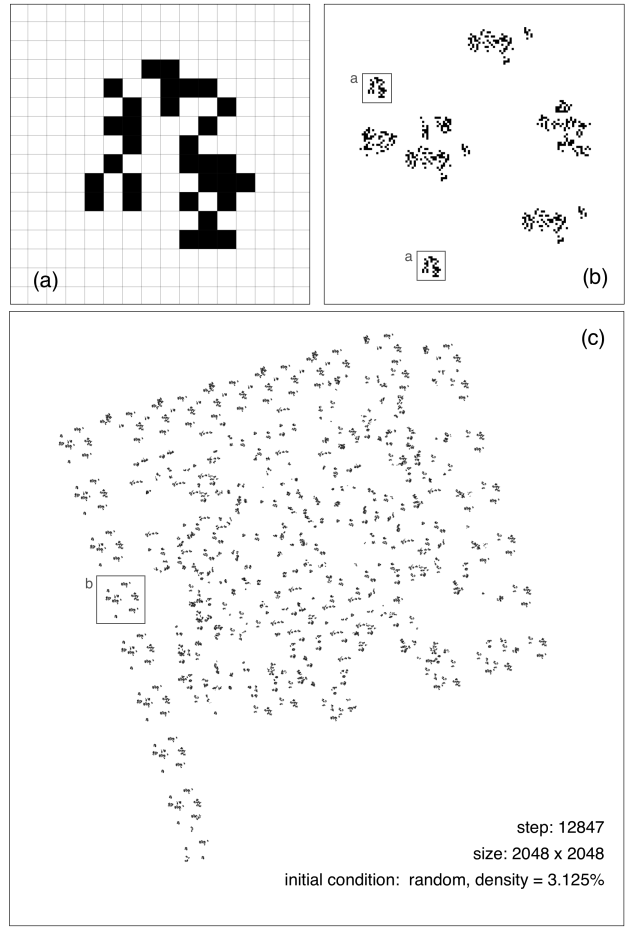

Figure 2: Sample outcome from the Outlier rule starting with a sparse random initial condition. (a) Two clusters on the smallest scale; (b) A self-replicating formation, assembled from a few clusters; (c) On the largest scale, an expanding complex with a semi-chaotic interior, bordered by replicating formations.

prompting the semi-chaotic phase to ensue.

A logical next step following this discovery would be to identify rules that support not only emergent replication, but also adaptation and structural evolution. While these capabilities have already been showcased with specially designed states (Sayama, 1999), the prospect of such rules existing within simpler CA remains uncertain. The Outlier has also demonstrated that emergent clusters on binary CA produced by similar rules can serve as building blocks for larger, complex constructs. This opens the path for potential manipulation of these clusters, acting as components of composite computing substrates, such as layered multiple binary CA. Since binary CA can be a superiorly efficient computing substrate, they have good potential to be part of large-scale complexity simulations, a prospect of interest to both OEE and machine intelligence (Ha and Tang, 2022).

References

Bar-Yam, Y. (2003). *Dynamics of Complex Systems*. Westview Press.

Bedau, M. A., McCaskill, J. S., Packard, N. H., Rasmussen, S., Adami, C., Green, D. G., Ikegami, T., Kaneko, K., and Ray, T. S. (2000). Open problems in artificial life. *Artificial life*, 6(4):363–376.

Chou, H. H. and Reggia, J. A. (1997). Emergence of self-replicating structures in a cellular automata space. *Physica D: Nonlinear Phenomena*, 110:252–276.

Ha, D. and Tang, Y. (2022). Collective intelligence for deep learning: A survey of recent developments. *Collective Intelligence*, 1(1).

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

Lehman, J. and Stanley, K. O. (2008). Exploiting open-endedness to solve problems through the search for novelty. In *Proceedings of the Eleventh International Conference on the Synthesis and Simulation of Living Systems, ALIFE 2008*, pages 329–336. MIT Press.

Sayama, H. (1999). A new structurally dissolvable self-reproducing loop evolving in a simple cellular automata space. *Artificial Life*, 5:343–365.

von Neumann, J. (1966). *Theory of self-reproducing automata*. University of Illinois Press.