

Cultivating the Garden of Eden

Randall D. Beer

Cognitive Science Program

Informatics Department

Luddy School of Informatics, Computing, and Engineering

Indiana University

rdbeer@indiana.edu

Garden of Eden (GOE) states in cellular automata are grid configurations that have no precursors; that is, they can only occur as initial conditions. Finding individual configurations that minimize or maximize some criterion of interest (e.g., grid size, density, etc.) has been a popular sport in recreational mathematics, but systematic studies of the set of GOEs for a cellular automaton have been rare. This paper presents the current results of an ongoing computational study of GOE configurations in Conway's Game of Life (GoL) cellular automaton. Specifically, we describe the current status of a map of the layout of GOEs and non-GOE in 1-density/size space, characterize how the density-dependent structure of the number of precursors varies with increasing grid size as we approach the point where GOEs begin to occur, provide a catalog of all known GOE configurations up to a grid size of 11×11 , and initiate a study of the structure of the network of constraints that characterize GOE versus non-GOE configurations.

1. Introduction

The Game of Life (GoL) is a well-known cellular automaton that was invented by John Conway and originally popularized by Martin Gardner in the pages of *Scientific American* in 1970 [1–4]. In the GoL, each cell in a square two-dimensional grid can be either ON or OFF. A cell changes state according to a simple rule based on its own state and the states of the eight cells immediately surrounding it. Continuing interest in this cellular automaton derives from the remarkable complexity that arises from the temporal evolution of this simple rule. Indeed, it is well known that the GoL is computationally universal.

Of the many interesting questions raised by the GoL, the Garden of Eden (GOE) problem has been one of the most enduring. A GOE configuration is one that has no precursors, that is, that can only occur as an initial state in the time evolution of a cellular automaton. The first GOE configuration in the GoL, a 33×9 pattern containing 226 ON cells, was discovered in 1971 by Roger Banks and verified by Don Woods [5, pp. 230, 248; 6, p. 49] Since that time, a small but devoted

group of enthusiasts has searched the space of GoL configurations, informally competing to find ever-smaller or less dense GOEs [7, 8].

A significant disadvantage of the focus on breaking records is that it tends to emphasize the characteristics of individual GOEs. This is unfortunate because GOEs are much more than just an entertaining challenge in recreational mathematics. Because GOEs serve as the basin boundaries of distinct global attractors, they play a central role in organizing the global dynamics of a cellular automaton. Accordingly, it would be useful to have a more systematic understanding of the properties of the ensemble of GOEs as a whole.

Work along these lines was recently initiated by Hartman et al. [9], who generated all square GOE configurations in the GoL respecting certain symmetries up to a grid size of 14×14 . They found that no GOEs exist below a grid size of 6×6 and that no GOEs satisfying certain symmetries exist below a grid size of 10×10 ; the question of whether or not nonsymmetric GOEs exist at intermediate grid sizes remains open. In addition, they found that the occurrence of GOEs was strongly dependent on the 1-density of configurations. Specifically, the number of symmetric GOEs was found to be normally distributed with respect to density.

The goal of this paper is to extend this analysis in four ways. First, we present the current status of a comprehensive map of the layout of GOEs and non-GOE in 1-density/size space. Second, we characterize how the density-dependent structure of the number of precursors varies with increasing grid size as we approach the point where GOEs begin to occur. Third, we provide and analyze a catalog of all known GOE configurations up to a grid size of 11×11 . Fourth, we initiate a study of the structure of the network of constraints that characterize GOE versus non-GOE configurations.

2. Preliminaries

Conway's GoL is a binary outer-totalistic cellular automaton on a two-dimensional rectangular grid [1–4]. The most common statement of the GoL update rules is:

1. Loneliness: any live cell with fewer than two live neighbors dies.
2. Survival: any live cell with two or three neighbors lives to the next generation.
3. Overpopulation: any live cell with more than three live neighbors dies.
4. Birth: any dead cell with exactly three live neighbors comes alive.

This set of rules can be stated more compactly as: “A cell will be ON in the next generation if either it is surrounded by three ON cells

or it is currently ON and surrounded by exactly two ON cells; otherwise that cell will be OFF,” which can be written as

$$\phi(s, \Sigma) = \delta_{\Sigma,3} + s\delta_{\Sigma,2}$$

where s is the current state of that cell, Σ is the number of its neighbors that are currently ON, and $\delta_{i,j}$ is the Kronecker delta function, which takes on the value 1 when $i = j$ and 0 otherwise. Alternate formulations also exist, such as [10]

$$\phi(s, \Sigma) = H(2 - (s + 2\Sigma - 6)^2)$$

where $H(\bullet)$ is the Heaviside function, which takes on the value 1 when its argument is greater than 0 and otherwise returns 0.

The global GoL evolution operator $f: \mathcal{G} \rightarrow \mathcal{G}$ is then obtained by applying ϕ elementwise to every cell in a grid $\mathcal{G}$. A configuration C of an $n \times n$ region of $\mathcal{G}$ is a specific assignment of 0s and 1s to the grid cells in that region. The $(n+2) \times (n+2)$ preimage of C , denoted $C' = f^{-1}(C)$, is defined as the set $(C' \mid f(C') = C)$. Due to the order-8 D_4 symmetry of the GoL, the same configuration can appear in multiple ways, depending on how it is embedded in the grid. We will refer to these different embeddings of the same configuration as *presentations*.

One of the simplest possible questions we can ask about the preimages of a configuration is, Are there any? This is the GOE problem for cellular automata. A configuration C is a GOE configuration if it has no precursors, that is, if $|f^{-1}(C)| = 0$. A (possibly nonsquare) minimal subset of a GOE configuration is called an *orphan*. The idea of a GOE configuration derives from the following theorem:

Theorem 1. The Garden of Eden. [11, 12]: For a d -dimensional cellular automaton $f: \Sigma^{\mathbb{Z}^d} \rightarrow \Sigma^{\mathbb{Z}^d}$ over a finite alphabet Σ , f pre-injective $\Leftrightarrow f$ surjective.

The content of Theorem 1 is most easily understood in its contrapositive form, $\neg (f \text{ pre-injective}) \Leftrightarrow \neg (f \text{ surjective})$, which can be written as

$$\exists_{c_1 \neq c_2} f(c_1) = f(c_2) \Leftrightarrow \exists_c |f^{-1}(c)| = 0.$$

In words, a function f fails to be pre-injective if there exist two distinct finite configurations of the same size that are mapped to the same configuration by f . Such pairs of configurations are called *twins*. A function f fails to be surjective if there exist configurations without precursors, that is, GOE configurations. Thus, the theorem states that a cellular automaton possesses GOEs if and only if it has twins. The applicability of this theorem to the GoL was noted very early by Alvy Ray Smith [5, pp. 230, 248; 6, p. 49]. A simple pair of twins in the

GoL are the 5×5 configuration consisting of all OFF cells and the 5×5 configuration consisting of a single ON cell at its center, both of which evolve to the quiescent state.

Because an $n \times n$ grid in the GoL has $2^{(n+2)^2}$ possible precursors, the problem of searching for GOEs is exponential in the grid size. A standard approach is to work backward from a given target configuration, performing a depth-first search through the possible inversions of each cell state filtered by local consistency requirements and backtracking whenever a conflict arises. For the GOE problem, the search can be terminated as soon as the first precursor is found. Hardouin-Duparc [13] improved on this method by deriving a finite automaton that accepted GOEs. More recently, GOE search has been formulated as a satisfiability problem and tackled with SAT solvers [9, 14, 15].

3. A Map of the Territory

The current status of our ongoing GOE search is shown in Figure 1 as a function of grid 1-density ρ and grid size n , along with the accompanying numerical data in Table 1. From previous work, it was known




Figure 1. Current state of knowledge regarding GOE configurations in the GoL for square grids of size n and density ρ . Regions where GOE configurations are known to exist are shown in blue. Regions where GOE configurations have been proven not to exist are shown in red, with all results above $n = 6$ new to this paper. The white gap between these two regions indicates the current frontier where the question is not yet settled. The numerical boundaries of these regions are given in Table 1. The density of the configuration with the minimum number of precursors for each n is marked with red points connected by solid red lines, with open red circles connected by dashed red lines indicating incomplete computations that are still in progress. Finally, the current mean density at which GOE configurations occur is marked with open blue circles connected by a dashed blue line.

Grid Size	GOEs	No GOEs	Unknown
1×1		(0, 1)	
2×2		(0, 4)	
3×3		(0, 9)	
4×4		(0, 16)	
5×5		(0, 25)	
6×6		(0, 36)	
7×7		(0, 49)	
8×8		(0, 11) and (49, 64)	(12, 48)
9×9		(0, 8) and (70, 81)	(9, 69)
10×10	(47, 68)	(0, 7) and (90, 100)	(8, 46) and (69, 89)
11×11	(45, 94) (excepting 47, 48, 51 and 91)	(0, 6) and (112, 121)	(7, 44) and (95, 111) and 47, 48, 51, 91

Table 1. A numerical summary of the areas shown in Figure 1. For each grid size considered in this paper, the range of numbers of ON cells is given for which GOEs are known to exist (blue in Figure 1), for which GOEs are known not to exist (red), and for which the status is currently unknown (white). Note that there are some possible exceptions for the GOE region of the 11×11 case where GOEs have not yet been found.

that no GOEs exist for grid sizes of 6×6 or smaller, that GOEs exist for grid sizes of 10×10 and higher, and that no GOEs satisfying certain symmetries exist at intermediate grid sizes [9]. For the results in this section, we employ a very efficient iterative implementation of backward search with all consistency-checking reduced to table lookups. In addition, we take full advantage of the D_4 symmetry of the GoL in order to reduce the number of unique configurations that must be examined. Finally, since our GOE searches are organized by 1-density, grid configurations of a given density are generated in coollex order, which has a very fast branch-free implementation [16]. Using this approach, we have extended previous results in three ways.

First, we have shown that no GOEs exist in 7×7 grids, a result that took several years to obtain given that there are 562 949 953 421 312 configurations to be examined.

Second, we have shown that no GOEs exist outside an n -dependent range of densities for $8 \leq n \leq 11$ (red region). Since the number of $n \times n$ grids containing k 1s is $\binom{n^2}{k}$, this region represents the elimination of an additional 271 687 723 328 802 configurations, for a total of 834 706 429 847 348 lying in the red region of the plot. This result

also took several years to obtain. Note that the number of possible configurations of an $n \times n$ grid containing between l and m ON cells is given by

$$\sum_{i=l}^m \binom{n}{i} = \binom{n}{l} {}_2F_1(1, l-n; l+1; -1) - \binom{n}{m+1} {}_2F_1(1, m-n+1; m-n; -1)$$

where ${}_2F_1$ is the ordinary hypergeometric function.

Third, we have extended the number of GOEs that are known to exist and the densities over which they exist for $n = 10, 11$ (blue region), providing a large catalog of GOEs for analysis. These GOE configurations will be examined in more detail in Sections 5 and 6. Note in Table 1 that there are a few possible exceptions for the 11×11 case.

The white region in Figure 1 indicates the grid sizes and densities that remain to be systematically examined, representing a total of 10^{83} configurations. Given the amount of computation that would be required to complete this map, it is worth exploring other perspectives that might provide important insights into the nature and layout of GOEs in this space.

4. Approaching the Garden

Our first strategy is to embed the GOE problem in a more general problem. If a GOE is a configuration with no precursors, then perhaps we should examine how the number of precursors that a configuration can have varies with grid size and 1-density. In particular, we will focus on those configurations with the smallest number of precursors as we consider larger and larger grids. Unfortunately, because we cannot know a priori which configurations will have the smallest number of precursors, we must compute the number of precursors for all configurations of a given grid size and density in order to find the smallest. But even though precursor counting is more computationally expensive than GOE testing, it carries more structure than a simple binary classification.

Our basic object of study in this section is a *precursor count envelope* (Figure 2). For a given grid size, this is simply a log plot of the minimum and maximum number of precursors as a function of density. The mean number of precursors is also indicated. To date, precursor count envelopes for up to $n = 6$ have been completed using the incremental aggregation algorithm with all applicable optimizations [17]. The $n = 7$ envelope is currently in progress (2^{49} total configurations; 71.67% completed).




Figure 2. Precursor count envelopes for grids up to size 7×7 . The upper and lower curves in each plot bound the log precursor count ranges for each density, while the orange central line denotes the mean log precursor count. Asterisks mark the location of global minima, whereas the arrows for the 6×6 case indicate the local maxima referred to in the main text. Non-log closeups of the minima are shown as insets for the 6×6 and 7×7 precursor envelopes. Density ranges for which the search is complete are shown in blue, while density ranges for which the search is ongoing are shown in red. The state of completion for the incomplete density region of the 7×7 precursor envelope is shown as a gray curve whose scale is indicated on the right.

Precursor count envelopes have a characteristic “triangular” shape, which becomes more defined as n increases. The leftmost vertex, which corresponds to the vacuum state (all 0s), always has the most precursors, whereas the rightmost vertex, which corresponds to all 1s, does not generally have the fewest precursors. The number of precursors of these two configurations grows as $O(2^{n^2})$ and $O(2^n)$, respectively [17]. Between these two extremes we see a relatively straight left edge broken by a single kink, a long diagonal right edge punctuated by small peaks, and an increasingly flat bottom edge with a single minimum peak. Note that the apparent flattening of the edges and peaks is at least partly due to the logarithmic vertical scaling of these plots.

Let us begin with the global minima along the bottom edges of the precursor count envelopes (indicated with asterisks in Figure 2). These represent the configuration(s) that are closest to being GOEs for each grid size because they have the smallest possible number of precursors for that size. These minimal configurations are shown in Figure 3(a), along with their precursor counts. A few interesting observations can be made. First, whereas the $n = 5$ minimal configuration clearly builds on the $n = 3$ case, the $n = 4, 6$ and 7 minimal configurations are each unique and become increasingly intricate as grid size increases. Second, the minimal configurations possess very different symmetries, with symmetry group orders of 1, 2 or 4. Third, the minimal precursor count grows with grid size up to $n = 6$, but then decreases for $n = 7$. We can see this pattern even more clearly if we plot maximal and minimal precursor counts for $n \leq 10$ (Figure 3(b)). This plot shows that while the maximum curve continues to grow with n , the minimum curve turns around after $n = 6$, eventually plunging to $-\infty$ in this log plot for 10×10 grids, which are known to contain GOEs. A final observation is that the densities at which the minimal configurations occur exhibit a meandering but generally decreasing trend as grid size increases (red lines in Figure 1).

Although not directly related to our study of GOE configurations, another feature of these precursor count envelopes is worth exploring. As mentioned earlier, the upper limits of precursor count envelopes exhibit a series of peaks at regular intervals (indicated by arrows for the 6×6 case in Figure 2). The number of peaks increases with grid size and they also become narrower and sharper. In addition, the largest peaks occur at intermediate densities and then decay in size as we move toward the extremes. A closer examination reveals that as the maximum precursor count decreases with increasing density, each “peak” actually consists of a flattening out of the curve followed by a sharp drop. Interestingly, the peaks themselves occur at densities that are specific fractions of n^2 . For example, for the 6×6 case, we see peaks at $\rho = 0, 1/6, 1/4, 1/3, 1/2, 2/3, 3/4, 5/6$ and 1.

Although the 7×7 envelope exhibits similar features, we can make no detailed quantitative comparison yet, since this envelope is still under construction.




Figure 3. The progression of minimal precursor counts (asterisks in Figure 2) with increasing grid size. (a) The configurations with the smallest number of precursors for each grid size up to 7×7 . ON cells are displayed in black and OFF cells are displayed in white. The number of precursors for each configuration is given in parentheses. The minimum precursor is always unique, but may have multiple presentations due to its symmetries. Note that the 7×7 (red) computation is still ongoing, so that it is possible that a configuration with a smaller number of precursors might still be found for this case. (b) Precursor count ranges for grids up to size 10×10 . The inset is a non-log plot of the minimum curve of the main log plot, more clearly showing the sharp turnaround at $n = 6$. Here blue curves indicate completed computations and red curves indicate ongoing computations. However, it is important to emphasize that even the red portions constitute upper bounds for the minimum curve; additional data can only lower this portion of the curve if configurations with fewer precursors are subsequently found.

What underlies these patterns in the maximum precursor counts for a given grid size? As Figure 4(a) illustrates for 6×6 grids, these peaks occur for configurations that are particularly regular, such as stripes or blocks of cells in the same state. As a general rule, contiguous empty regions have the largest number of precursors and contiguous filled regions have the smallest number. Perhaps the peak configurations represent a balance between these two effects. This gives the impression that as more ON cells are added, the precursor count asymptotes until some packing limit is reached, and then a reorganization occurs into a new pattern. An examination of a representative maximal configuration for each possible number of ON cells shows that this is indeed the case (Figure 4(a)). A similar sequence of

reorganizations is observed for minimal precursor count configurations, but here the patterns are much more complicated (Figure 4(b)).




Figure 4. A representative example of each extremal precursor count configuration in 6×6 grids as the number of ON cells varies from 0 to 36. ON cells are displayed in black and OFF cells are displayed in white. (a) Maximal precursor count configurations. The configurations corresponding to the maximal precursor count peaks in the 6×6 envelope (denoted by arrows in Figure 2) are outlined in red along with their corresponding densities. (b) Minimal precursor count configurations. The globally minimal configuration is outlined in red.

5. Cataloging the Garden

Our second strategy is to expand the set of known 10×10 and 11×11 GOEs as much as possible in order to examine their collective properties. We will start by reconstructing the symmetric GOEs identified by Hartmann, Heule, Kwekkeboom and Noels (henceforth HHKN) [9] and then build from there. We will make use of the only two symmetry operators defined by HHKN for a square grid $\mathcal{G}$ that produce 10×10 and 11×11 GOEs:

$$\sigma_{90}(90 \text{ degree rotation}): \mathcal{G}_{i,j} = \mathcal{G}_{j,n+1-i} \text{ for } i, j \in (1, \dots, n)$$

$$(\sigma_D, \sigma_{180})(\text{combined diagonal reflection and } 180 \text{ degree rotation}):$$

$$\mathcal{G}_{i,j} = \mathcal{G}_{j,i} = \mathcal{G}_{n+1-i,n+1-j} \text{ for } i, j \in (1, \dots, n)$$

For the 10×10 case, we found eight GOEs satisfying σ_{90} and 62 GOEs satisfying (σ_D, σ_{180}) , for a total of 70. These are identical to the values reported by HHKN. We then generated all 1-, 2-, 3-, 4- and 5-bit perturbations from this set of symmetric GOEs, obtaining

456, 1740, 4864, 10 384 and 17 036 new GOEs, respectively. Finally, we discovered an additional 1000 GOEs through less systematic random sampling, for a grand total of 35 550. The density histogram of this sample of 10×10 GOEs is shown in Figure 5(a). It is roughly normal with a mean of 0.588 and a standard deviation of 0.0295 but exhibits significant fluctuations. The minimum number of ON cells in this set of GOEs is 47 (Figure 5(b), left) and the maximum is 68 (Figure 5(b), right). Removing D4-symmetric redundancy from our full set of 10×10 GOEs reduces their total to 4617.




Figure 5. Properties of our set of 35 550 10×10 GOEs. (a) The density histogram for this set, with mean 0.588 and standard deviation 0.0295. (b) Presentations of the 10×10 GOE configuration with the least (47; left) and most (68; right) ON cells. (c) A sample of some of the other symmetric GOEs from our set.

For the 11×11 case, we found 41 702 GOEs satisfying σ_{90} . This is 342 fewer than the 42 044 reported by HHKN. However, it is interesting to note that removing D4-symmetric redundancy from our set leaves us with 21 022 GOEs, which is exactly half of the number

HHKN reported. We then generated all 1-bit perturbations of this set of symmetric GOEs, obtaining 565 280 new GOEs for a grand total of 606 982. The density histogram of this sample of 11×11 GOEs is shown in Figure 6(a). It is roughly normal with a mean of 0.588 and a standard deviation of 0.0025 but exhibits significant fluctuations. The minimum number of ON cells in this set of GOEs is 45 (Figure 6(b), left) and the maximum is 94 (Figure 6(b), right). Removing D4-symmetric redundancy from our full set of 11×11 GOEs reduces their total to 91 900.




Figure 6. Properties of our set of 606 982 11×11 GOEs. (a) The density histogram for this set, with mean 0.588 and standard deviation 0.0025. (b) Presentations of the 11×11 GOE configuration(s) with the least (45; left) and most (94; right) ON cells. (c) A sample of some of the other symmetric GOEs from our set.

What are the prospects for finding any 8×8 or 9×9 GOE configurations? HHKN showed that no symmetric 8×8 or 9×9 GOEs exist, but the nonsymmetric cases remain open. The only evidence we have for this question comes from Figure 1, where we observe that the (currently incomplete) boundaries of the blue GOE region, as well as its current mean GOE 1-densities, all seem to be converging to a point around $n = 9$. If these boundaries are close to the final ones, and if they can be linearly extrapolated, then this suggests that there will be

no 8×8 GOEs, but that 9×9 GOEs might exist in a narrow region around a density of $\rho = 0.59$ (corresponding to about 47 ON cells). Interestingly, the 9×9 configuration with the fewest precursors (1394) that we have so far observed contains 49 ON cells.

6. The Anatomy of Flowers of Eden

Our third and final strategy is to begin to examine the structure of the networks of constraints that underlie GOEs. A first step in this direction is to consider the following *dilation decomposition* of a GOE. Beginning with the central set of cells in an $n \times n$ GOE (a single cell if n is odd and a 2×2 grid if n is even), we count the number of precursors that each subsequent increase in window size by $n + 2$ has as we incrementally dilate our window into the GOE until its full extent is reached (which has zero precursors by definition). Our goal is to examine the scaling of the precursor count as we approach the GOE. An example of this decomposition for a 10×10 GOE is shown in Figure 7(a). Figures 7(b) and 7(c) compare this scaling for GOEs (blue) and density-matched non-GOE (red) in 10×10 and 11×11 grids, respectively.




Figure 7. Precursor scaling across dilation decompositions of GOE (blue) and density-matched non-GOE (red) configurations. (a) A dilation decomposition of the 10×10 GOE on the right, with the number of precursors shown below each step of the decomposition. (b) Precursor scaling across dilation decompositions of our D4-symmetry-reduced set of 4617 10×10 GOEs (blue) compared to the same number of density-matched non-GOEs (red), with individual curves shown with lighter thin lines and means shown with darker thick lines. (c) Precursor scaling across dilation decompositions of a random sample of 5000 configurations from our D4-symmetry-reduced set of 91 900 11×11 GOEs compared to the same number of density-matched non-GEOs.

Several interesting observations can be made about these plots. For small window sizes, GOEs and non-GOE are indistinguishable. In fact, on average, the central cores of GOEs actually have more precursors than those of non-GOE in the 10×10 case. Initially, both sets of curves also grow exponentially with increasing window size. However, because the rates of growth are smaller for the GOE curves, these curves quickly approach the lower limit of the range of precursor counts. Eventually, the GOE curves turn around completely, actually decreasing with increasing window size before plunging to $-\infty$ in this log plot. Thus, we can clearly see the interplay between two factors: the exponential growth in precursor count with increasing grid size and the exponential decay of precursor count with increasing constraints. Although the network of possible constraints on precursors is initially outpaced by the exponential growth of the number of possible precursors with grid size, it eventually tightens sufficiently for $n \geq 6$ for GOEs to occur.

Note that the increasingly wide precursor count separation between the dilation decomposition curves for GOEs and non-GOE does not mean that near misses do not occur. For example, if we flip just one bit of the 10×10 GOE shown in Figure 7(a), we can find non-GOE with as few as 14 precursors. Also note that analogous precursor count scaling curves can be constructed for any other incremental approach to a GOE, such as the greedy spiral method that Nicolay Beluchenko has used to construct a GOE [18].

These results highlight the special nature of the network of constraints that underlie GOEs. It would thus be very interesting to attempt to characterize such networks in more detail. While such a project lies outside the scope of the current paper, it is worth briefly visualizing how it might work. For example, consider the precursors of the {ON, OFF} configuration consisting of an ON cell adjacent to an OFF cell (Figure 8). Imagine stacking the 140 possible 3×3 precursors for an ON cell and the 372 possible precursors of an OFF cell (Figure 8, left). This generates up to $140 \times 372 = 52\,824$ possible precursor configurations for the joint {ON, OFF} pattern.

However, because these two cells are adjacent in the grid, the right two columns of a potential ON precursor and the left two columns of a potential OFF precursor (shown in red) must match for any pairing that forms a valid precursor to the joint configuration (Figure 8, right). This determines a constraint network (red lines), or, equivalently, a consistency network (green lines) that strongly restricts the allowable pairings. If we incrementally add one component at a time to analyze larger and larger patterns, then these networks will always be (bipartite) graphs. However, if we consider adding multiple components of a target pattern simultaneously, then both the constraint and consistency relations in general define hypergraphs. In

either case, in this formulation, GOEs correspond to networks whose consistency set is empty. Note that this representation makes explicit both the growth of possible pairings due to the multiplicative character of their combination and the decay of allowable pairings due to the increasing numbers of constraints that are brought to bear on these combinations.




Figure 8. Visualization of the constraint (red) or, equivalently, consistency (green) network for the precursors of a simple 2×1 configuration consisting of an ON adjacent to an OFF cell. (Left) The possible 3×3 precursors for each target cell individually, with the overlapping portions that constrain their combination highlighted in red. (Right) An illustration of the final set of 4×3 precursors of the joint {ON, OFF} target configuration that satisfy the consistency/constraint network.

7. Conclusion

In this paper, we have extended Hartman et al.'s [9] systematic study of Garden of Eden (GOE) configurations in the Game of Life (GoL) in four ways. First, we have constructed a map of the layout of the currently known GOEs and non-GOE in 1-density/size space. The most notable result here was proving that no 7×7 GOE configurations exist, but we also tightened the bounds over which 8×8 and 9×9 GOEs might exist as well as the bounds over which 10×10 and 11×11 GOEs can exist. Second, we have used precursor count envelopes to characterize how the density-dependent structure of the number of precursors varies with increasing grid size as we approach the point where GOEs begin to occur. The significant result here was showing that the minimal precursor count, which initially grows with grid size, turns around and begins to shrink after $n = 6$, but we also identified the minimal precursor count configurations up to 7×7 and studied features of the maximal precursor count configurations as well. Third, we expanded the set of known GOEs up to a grid size of 11×11 . Finally, we initiated a study of the internal structure of GOEs

by introducing dilation decomposition of GOEs, studying the precursor count scaling of such decompositions, and sketching a way to visualize the networks of constraints that underlie GOEs.

There are a number of directions for future work. First, the analysis in this paper is still woefully incomplete. The density/size map, the precursor count envelopes and the set of known GOEs all require significant expansion, with a special focus on the 8×8 and 9×9 grids with densities in the range of 0.55–0.65 where their minimal precursor count configurations seem to reside. This will obviously require significantly more computational resources. In the long run, SAT approaches may be better for GOE searches, but something like the incremental aggregation algorithm may still be well-suited for precursor counting. Second, it would be interesting to extend the analysis to nonsquare grids, with the current density/size map representing merely a slice through this now three-dimensional space. Third, it would be useful for the community to set up and maintain a public database of known GOE configurations, ideally with an associated orphan database. Finally, and perhaps most fundamentally, a serious study of the structure of the networks of constraints underlying GOE versus non-GOE configurations needs to be undertaken.

Acknowledgments

Computation performed for this research was supported in part by Lilly Endowment, Inc., through its support for the Indiana University Pervasive Technology Institute, and in part by the Indiana METACyt Initiative. The Indiana METACyt Initiative at IU is also supported in part by Lilly Endowment, Inc.

References

- [1] M. Gardner, “Mathematical Games: The Fantastic Combinations of John H. Conway’s New Solitaire Game ‘Life’,” *Scientific American*, 223(4), 1970 pp. 120–123. www.jstor.org/stable/24927642.
- [2] E. R. Berlekamp, J. H. Conway and R. K. Guy, *Winning Ways for Your Mathematical Plays*, New York: Academic Press, 1982.
- [3] A. Adamatzky, ed., *Game of Life Cellular Automata*, New York: Springer, 2010.
- [4] N. Johnston and D. Greene, *Conway’s Game of Life: Mathematics and Construction*, Lulu.com, 2021. (Feb 1, 2023) conwaylife.com/book.
- [5] M. Gardner, *Wheels, Life, and Other Mathematical Amusements*, New York: W. H. Freeman, 1983.

- [6] W. Poundstone, *The Recursive Universe: Cosmic Complexity and the Limits of Scientific Knowledge*, Chicago: Contemporary Books, 1985.
- [7] LifeWiki. “Garden of Eden.” (Feb 1, 2023) conwaylife.com/wiki/Garden_of_Eden.
- [8] A. Flammenkamp. “Garden of Eden / Orphan.” (Feb 1, 2023) www.homes.uni-bielefeld.de/achim/orphan.html.
- [9] C. Hartman, M. J. H. Heule, K. Kwekkeboom and A. Noels, “Symmetry in Gardens of Eden,” *The Electronic Journal of Combinatorics*, 20(3), 2013 P16. doi:10.37236/2611.
- [10] S. Adachi, F. Peper and J. Lee, “The Game of Life at Finite Temperature,” *Physica D: Nonlinear Phenomena*, 198(3), 2004 pp. 182–196. doi:10.1016/j.physd.2004.04.010.
- [11] E. F. Moore, “Machine Models of Self-Reproduction,” *Proceedings of Symposia in Applied Mathematics*, 14, 1962 pp. 17–33.
- [12] J. Myhill, “The Converse of Moore’s Garden-of-Eden Theorem,” *Proceedings of the American Mathematical Society*, 14, 1963 pp. 685–686. doi:10.1090/S0002-9939-1963-0155764-9.
- [13] J. Hardouin-Duparc, “Paradis terrestre dans l’automate cellulaire de Conway,” *Revue Française d’Automatique Informatique et Recherche Opérationnelle. Informatique Théorique*, 8(3), 1974 pp. 63–71. doi:10.1051/ita/197408R300631.
- [14] S. Bain, “Time-Reversal in Conway’s *Life* as SAT,” in *Proceedings of the 20th Australian Joint Conference on Advances in Artificial Intelligence (AI 2007)*, Gold Coast, Australia (M. A. Orgun and J. Thornton, eds.), Berlin, Heidelberg: Springer-Verlag, 2007 pp. 614–618.
- [15] D. Knuth, *The Art of Computer Programming: Satisfiability, Volume 4, Fascicle 6*, Addison-Wesley, 2015.
- [16] F. Ruskey and A. Williams, “The Coolest Way to Generate Combinations,” *Discrete Mathematics*, 309(17), 2009 pp. 5305–5320. doi:10.1016/j.disc.2007.11.048.
- [17] R. D. Beer, “Computing Aggregate Properties of Preimages for 2D Cellular Automata,” *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 27(11), 2017 111104. doi:10.1063/1.5006143.
- [18] N. Belcenko. “The On-Line Encyclopedia of Integer Sequences.” (Feb 1, 2023) oeis.org/A196447.