Homeodynamics in the Game of Life

Keisuke Suzuki and Takashi Ikegami

General Systems Sciences,
The Graduate School of Arts and Sciences,
The University of Tokyo,
3-8-1 Komaba, Tokyo, 153-8902, Japan

Abstract

In this paper we study the emergence of homeodynamics and adaptation in a two-layer system of the Game of Life in which the Game of Life in the first layer couples with another cellular automata system in the second layer. Homeodynamics is defined here as a space-time dynamic that regulates the number of cells in state-1 in the Game of Life layer. A genetic algorithm is used here to evolve the rules of the second layer to control the pattern of the Game of Life. We discovered that there are two antagonistic attractors that control the numbers of cells in state-1 in the first layer. The homeodynamics sustained by these attractors are compared with the homeostatic dynamics observed in Daisy world.

Introduction

Living systems require a stable and sustainable structure on top of unstable and highly chaotic open environments. The maintenance of such a structure is called "homeostasis", as named by Cannon (1932), and became one of the central themes in Cybernetic studies(Wiener, 1948). Several mechanisms underlying homeostasis have been proposed and they have become a guiding principle of our everyday technology. For example, positive/negative feedback loops and afferent/efferent copies are well studied and developed.

The study of homeostasis has revealed those mechanisms, but they are often introduced as a controlling device and the evolution of homeostasis itself has not been discussed seriously. People continue to study ecological homeostasis, in particular after Lovelock (1972) proposed his Gaia hypothesis. The Gaia hypothesis posits that the complex and global network of living/nonliving systems we observe self-organizes into homeostatic states. The Gaia hypothesis has been theoretically examined by Watson and Lovelock (1983) by developing the Daisy world model, a simple implementation of the Gaia theory. In Daisy world, temperature should be sustained at a certain range independent of the environmental temperature. Harvey (2004) calls the mechanism underlying the Daisy world a "rein control," a controlling mechanism which serves to pull the temperature toward the viability zone.

What has been missing thus far in the study of Daisy world is the self-organizing and dynamic nature of homeostasis. Ikegami and Suzuki (2008) studied a dynamic version of Daisy world controlled by spatio-temporal chaos. Because the homeostasis here is dynamically sustained, we refer to this as homeodynamics. Moreover, Homeodynamics doesn't simply hold the average temperature constant, as in a conventional Daisy world simulation, but instead aims to keep the temperature variation around the average. Holding variation brings adaptability into the Homeodynamic system, as it can respond to novel environmental conditions. This is the most significant characteristic of homeodynamics, which we will also focus on in this study.

With respect to this adaptability of homeo-systems, Ashby (1960) proposed an interesting design principle for the brain and for life forms as a whole which was mainly driven by homeostasis. He posited that the adaptive behavior of life is only an outcome of homeostatic properties and proposed a different type of homeostatic system called an ultra-stable system. This new system has two feedback loops. The primary feedback loop is driven by a mutual interaction between an organism's complex sensory and motor channels and the environment. Another feedback loop develops from the interaction between viability constraints and the relevant reacting parts via the essential variables that control the reacting parts. Usually, the second feedback loop is intended to change the meta-parameters of the system. When parameter values are outside of the viability constraints, the second feedback loop adjusts the essential parameters to let the system move towards a more stable state. These characteristics of the ultra-stable system share common features with our homeodynamic systems. That is, two dynamics co-exist in the same system with different time scales and they cooperatively control the homeostasis by keeping sufficient fluctuations in the system. In other words, we need both stable and unstable dynamics to develop homeostasis and adaptation at the same time.

In this paper, we study the notion of homeodynamics and adaptation by using Conway's Game of Life. A major drawback of most homeostatic models, including ours, is that

many systems can be too stable in the sense that they can survive without paying significant costs (or in the other word a system never dies). Therefore our challenge is to see how homeostasis can emerge even in a very unstable world, as in the Game of Life. A second objective of the paper is to see how robust homeostatic behavior is balanced with purposeful behavior such as memorizing the initial states. Robust homeostasis (keeping the system's state density constant) can be achieved by making a system insensitive to the initial Life density. But memorizing the initial state density means that a system should become sensitive to the initial Life patterns. These two opposite properties must be balanced within the same system.

In the next section we describe how to use the Game of Life to study homeostasis. In §3, we describe and analyze the results, and in §4, we discuss the observed characteristics of homeodynamics and adaptation in the Game of Life and also attempt to apply these results to the more general notions of homeodynamics.

The Model

The basic idea of the model is inspired from work by Taylor (2004). In his model, the system under examination consists of two layers of cellular automata: one is the Game of Life, and the other layer serves to control the Game of Life pattern. The dominating cell layer does not have to be governed by the rules of Life, but instead can be driven by a different rule set. Taylor evolved the rules of this layer by using an evolutionary algorithm to control a virtual sensorimotor flow arranged on the first layer. For example, when an input bit on the first layer is state 1, the output (target) area should have many state-1 cells. This contingency between input and output bits is mediated by the intermediate area, in which some of the bits are governed by the rules in the second layer.

The purpose of Tayler's study is to examine an unseparated body-environment boundary and to see the emergence of the boundary itself. We will also investigate this point, but here we will use his approach to study homeostasis. Our setup is described below.

The model consists of a 2D cellular automata running the Game of Life. Extra rules encoded in the genome can override Life states in a certain part of the CA space. A *target area* is also designated, which is significant for the central tasks of the simulation.

In Taylor's work there are two kinds of rules for genes: conditional and temporal genes. The conditional gene is activated when a certain requirement of neighboring cells is satisfied. The temporal gene is activated when a certain time step passes. In both cases, the genes have a coordinate specifying the target cell.

Here we used only the conditional rule. We modified the rule to become a so-called "totalistic rule" in which the gene only takes into account the number of neighboring state-1 cells, and not any specific neighboring patterns. Figure 1 illustrates the gene components we used. Note that in our model, all the genes are activated every time step, due to the construction of the genome and the use of this totalistic rule. The gene consists of a 26-bit binary string, 18 bits of which represent a totalistic rule. As is usually the case, the 2D CA rule is specified by the pair of numbers Bx/Sy, which specifies when to change the cell's state to 1, whether its own state is 1 or 0, respectively. For example, the Game of Life is represented by B3/S23. We use 9 bits to represent those two parameters.

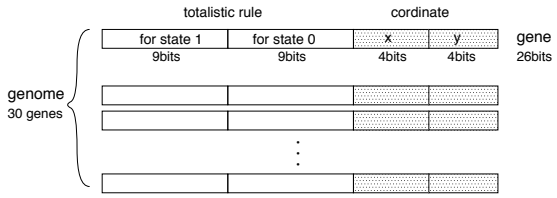

Figure 1: A gene component consists of a 26-bit binary string, 18 bits of which encode the totalistic CA rules. The remaining bits encode the coordinate of the site where the rule activates. A group of these genes is called the genome. The simulation is run using Game of Life dynamics and one genome.

The remaining 8 bits encode the spatial position of the site controlled by the second layer (4 bits for each x and y coordinate). The length of each gene is fixed and each gene specifies a particular single site in the 16×16 cell space.

The total cell space is given as a square of the size 40×40 and the intermediate area controlled by those genes is given as a square of the size 16×16. The target area is defined as a square space and all three squares share a common center. The size of the target area is 32×32 bits and includes the intermediate area (See Fig. 2). The boundary of the space is always set to state-0.

The target behavior of the model

While evolving the rule set of the intermediate cell space we will study homeostatic behaviors observed in the Game of Life. Instead of a temperature value as in Daisy world, we use the density of cells in the state-1 as the target variable to keep constant.

In Daisy world, the system consists of two types of flowers, black and white daisies, with their local temperature values. The black daisies increase the temperature and the white daisies decrease the temperature. By setting the growth rate of these flowers according to the local temperature, both flowers show positive feedback effects in their relationship to changing local temperatures. While the black daisies increase both the temperature and their population, the white daisies decrease the temperature and increase their population. Loosely linked by the the two local tempera-

Figure 2: The space used for the simulation. Filled squares represent the area governed by Game of Life rules. Lined squares represent the area which the genome specifies. The inner square shows the area in which genes override the Life states, and the outer square depicts the target area. The dynamics is only evaluated in this target area.

ture, the global temperature is sustained constantly while both population of daisies change according to environmental temperature changes.

This result shows that the homeostatic behavior does not result from insensitivity to the environmental stimulus, but is actively achieved by an adaptive coupling of components which are sensitive to the environment.

In order to observe the underlying dynamics of homeostasis in the Game of Life, we constructed three different tasks.

task A Sustain the same density of the target area regardless of the initial Life pattern density.

task B Control the density of the target area, making it proportional to the given initial Life pattern density.

task C Control the density of the target area, making it inversely proportional to the given initial Life pattern density.

The first task is designed to directly aid the development of a homeostatic behavior in the Game of Life patterns. We will see how this behavior is achieved in the Game of Life space. The second and third tasks are intended to facilitate the development of sensitivity to the given conditions. These target behaviors in task B and C are not directly connected to homeostatic behavior, but might be connected to the property of adaptability in homeostasis.

GA

In order to observe those behaviors, the CA rules encoded in the genes are evolved by a simple genetic algorithm (GA). We prepared 30 genomes in a population, each of which consists of 30 genes, which specify each spatial location and the rule content of the 30 CA rules in the intermediate area. Mutation occurs at every site of the gene at a certain rate per bit, which modifies the spatial locations and the rule content. The number of genes also changes during this process, so the number of CA sites in the intermediate area also varies. For the selection algorithm, we used a roulette selection procedure and chose an elite strategy.

Parameter	Values
population	30
the mutation rate	0.05
the crossover rate	0.01
the mutation rate for genome length	0.01
the number of elite	5
initial density(higher)	0.5
initial density(lower)	0.0
evaluated duration(time steps)	500

Table 1: The parameters used in our simulation

The GA goes through both an evolving phase and a testing phase. In the evolving phase, 30 genomes are evolved as a unit against two different initial states with lower and higher density patterns. Here the lower density is set to 0 and the higher density is set to 0.5.

Fitness of the genome is calculated by finding how the average density of state-1 within the target area compares to the specified *target densities*. The target densities in the three tasks are set as follows.

$$\begin{aligned} 0.5 \quad & (\text{for task A}) \\ d_I \quad & (\text{for task B}) \\ 1 - d_I \quad & (\text{for task C}) \end{aligned}$$

Here, d_I is the intial density of state-1 cells, set as either 0 or 0.5.

Note that we only use a fixed random initial pattern for all evolutionary processes. Once the system evolves, it develops a sufficient generalization capability; the system can do well with new initial patterns. However, a full generalization capability is difficult to obtain. We will revisit this point later. Table 1 shows the parameter values used in this experiment.

Result

Evolved Dynamics

In each of three tasks, the genomes in our population were evolved for higher fitness. Figure 3 show the temporal changes of the state-1 density of the fittest genome in each task.

Two lines on the graphs are shown, one for the low initial density case (null pattern) and another for the high density case (0.5). These density values were used in the GA dynamics.

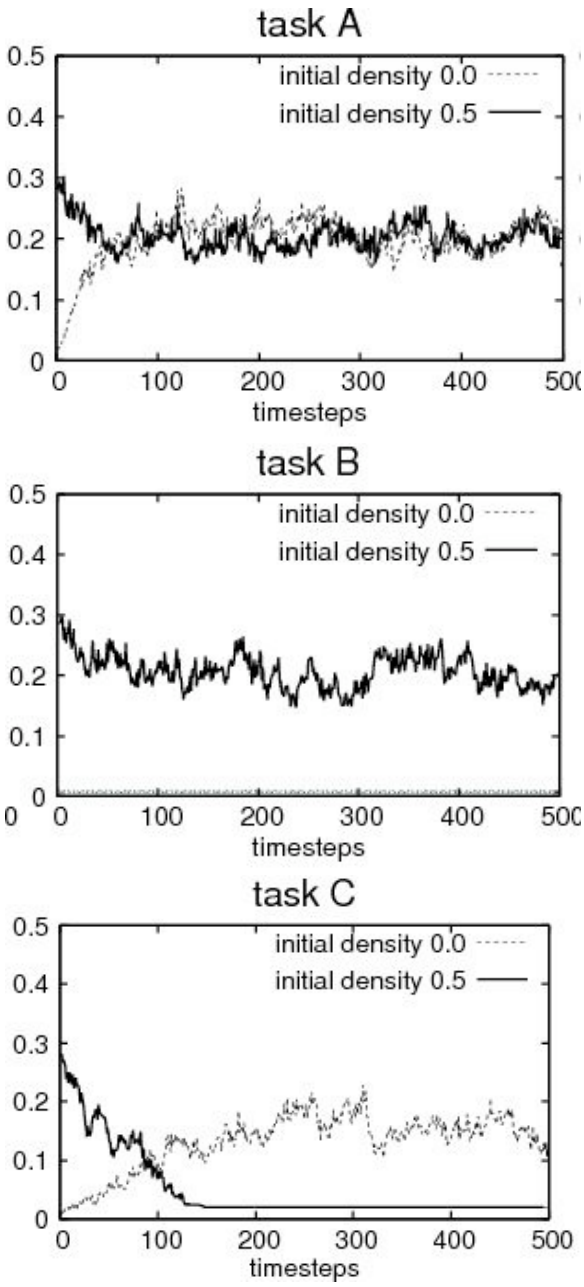

Figure 3: Temporal changes of the state-1 density of the fittest genome in each task. The initial Life patterns used here are the same as those used during the GA procedure. When begun with both low and high density initial states, the state-1 densities are maintained at around 0.2(task A). With the higher-density initial state, the state-1 density is kept around 0.2, but with the lower-density initial state, the state-1 density decreases to 0(task B). The result observed here is the inverse of task B. The higher-density initial state shows a state-1 density which drops almost to 0, but the lower-density initial state causes a growth in state-1 density until the density reaches approximately 0.15(task C).

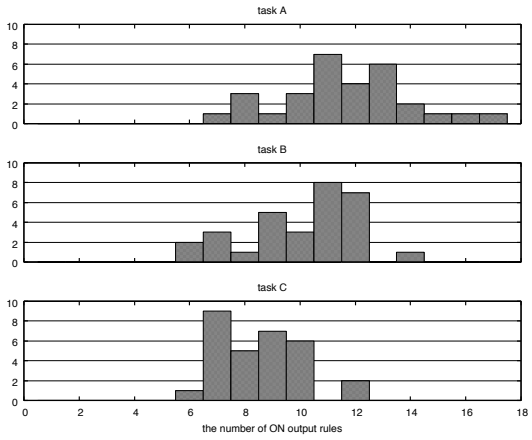

Figure 4: Histograms of the number of state-1 outputs out of the overall outputs for the fittest genomes in each of the three tasks. Task A and B genomes have a biased distribution towards larger values, while task C is biased in the opposite direction.

In task A, the density almost always reaches the same value of approximately 0.2 regardless of the initial density. We can see here that this attracting state is maintained by the generators of the Game of Life pattern, which will be discussed later. Cloud-like patterns in the Life space are generated by the evolved CAs.

In task B, the initial low-density state almost always creates a sparse pattern in the target area. Using the higher initial density state, the average resultant density state tends to fluctuate around a value of 0.2. The genomes generally increase the density if the target area is surrounded by a high-density pattern; the density tends to decrease in the sparse case. A generator of this type creates cloud-like patterns in the higher-density environment, but suppresses cloud-like patterns in the lower-density environment.

In task C, the densities are altered to be inversely proportional to the initial states. When the initial state has a low density, the evolved CA creates a high density state. In contrast, when presented with the high-density initial state, the density is decreased until all Life patterns are diminished. In the higher-density initial state, we expect an increase in the resultant density. The evolved CAs have to inhibit the spontaneous generator to decrease the state-1 density. Thus, the genome in task C behaves like an activator in the initial low-density state, but behaves like an inhibitor in the initial high-density state.

CA rules of the second layer

Each of the 30 sites in the second layer has a different CA ruleset. One way to characterize them is to compute the number of state-1 outputs as a fraction of overall outputs (e.g. the Game of Life has 3/16). Figure 4 shows a histogram comparison of these state-1 outputs for the evolved

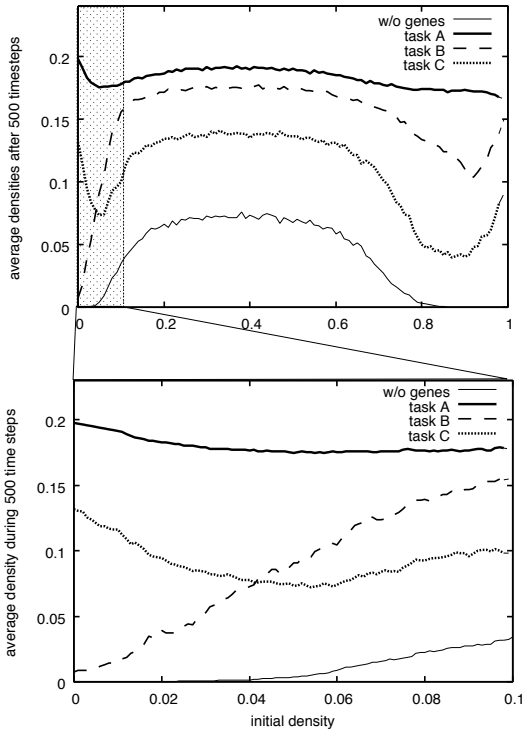

Figure 5: The average densities observed during 500 time steps when starting from different initial densities. The evolved genomes in task A, B, C and runs in which only Game of Life rules exist are compared. Each value is averaged over 100 different initial Life patterns.

genome sets.

Before evolving genome sets, the output is normally distributed around 8. The rulesets of tasks A and B are biased toward higher state-1 outputs, but those of task A are more biased towards larger values than those of task B. In task C, the rule sets are biased in the opposite direction.

In task A, the genome generates a cloud-like pattern regardless of the initial pattern. The generator is a strong activator of the state-1 LIFE pattern. Task B (proportion) also generates a cloud-like pattern, but only in the higher-density initial condition. The decrease of the density in the lower initial pattern may require a weaker activator in this task than that in task A. In task C (inverse proportion), the genome has to inhibit the cloud-like patterns in the higher-density initial state and activate the cloud patterns in the lower-density initial states.

The genomes in task B and C can be compared with the black and white daisies of Daisy world. The genome in task B has a tendency to output state-1, which can be regarded as similar to the black daisy which increases the local temperature. Likewise, the genome in task C has a tendency to suppress state-1, which can be regarded as similar to the white daisy which decreases the local temperature. How-

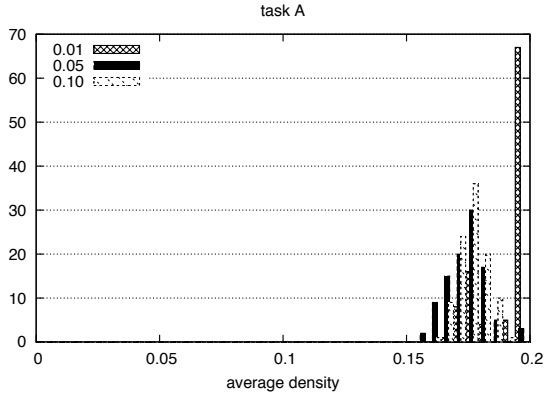

Figure 6: Histogram of the average densities observed from the task A genome in 100 samples with different initial Life patterns. Here, initial densities of 0.01, 0.05, and 0.10 are used.

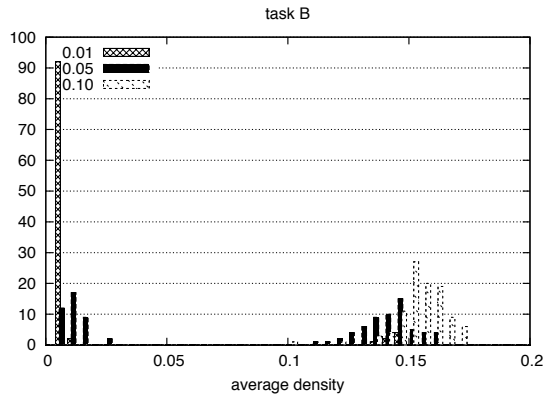

Figure 7: Histogram of the average densities observed from the task B genome in 100 samples with different initial Life patterns. Here, initial densities of 0.01, 0.05, and 0.10 are used.

ever, it should be noted that the behavior of the evolved CA rule is also affected by their spatial configuration. Thus, the state-1 output levels do not fully reflect how many cloud-like patterns these evolved CA rules can make by themselves.

Generalization

While the genomes under discussion here have been trained with only two fixed initial Life patterns, the final genome acquired a generalization capability to some extent. After training the genome set with these two different initial densities of 0.1 and 0.5, we have tested the evolved genome against other initial densities. Figure 5 shows the average density after 500 GA time steps for the given initial density. Each density is averaged over 100 different initial Life patterns of a given state-1 density. For a comparison, we also show the average density obtained with only the origi-

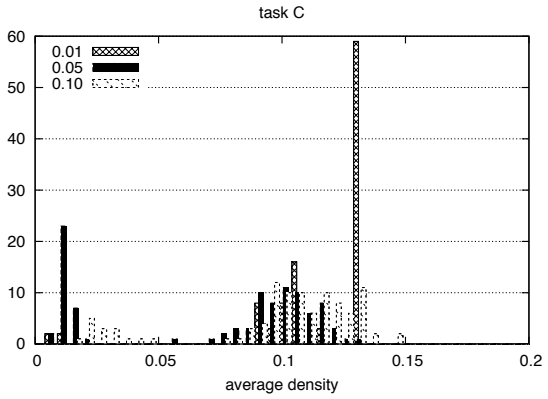

Figure 8: Histogram of the average densities observed from the task C genome in 100 samples with different initial Life patterns. Here, initial densities of 0.01, 0.05, and 0.10 are used.

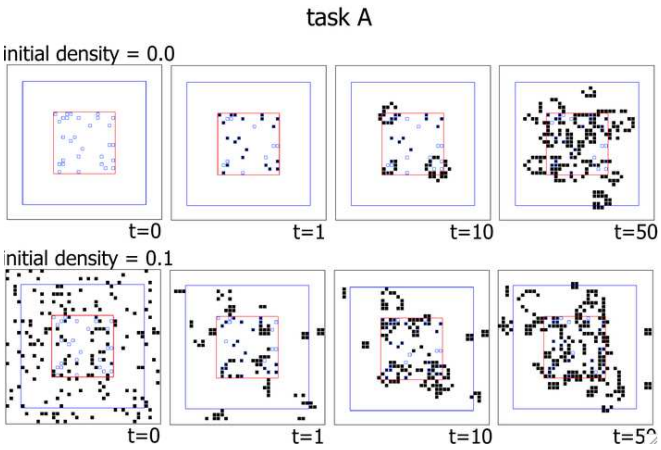

Figure 9: Snapshots of the Life dynamics of the fittest genome in task A. The top columns show the results when the initial density is zero, whereas in the bottom column the initial density is 0.1. In both cases, clouded Life patterns spread out from the genome area.

nal Life dynamics, without any second layer.

The genome evolved by task A achieves almost identical values of high average density against a wide range of initial densities from 0.0 to 1.0. Compared with the original Game of Life, this evolved genome sustains a higher density, particularly around the lowest and highest initial state. This is achieved by a state-1 generator which increases the state-1 density regardless of the initial states. However, similar behavior can be seen when one adds noise to the Game of Life. So the genome does not regulate state-1 density, but rather works as a "random generator".

We are not certain how many attractors this two-layered system has when starting from identical initial density states. Fig. 6 shows that there are two peaks in the histogram which correspond to different attractors in the system. Regardless of the final attractor, the genome evolved via task A has similar final density states, i.e. 0.15 -0.19.

Figure 9 shows snapshots of the fittest genome found in task A with the two different initial densities. In both cases, the Game of Life patterns become similar cloud-like shapes after time passes.

Tasks B and C also inherit the tendency to increase and decrease the state-1 density from the condition in which evolution occurs.

When the initial density is small(< 0.1), the genome evolved via task B shows a linear dependence on that initial density, while the genome evolved via task C shows the inverse dependency. The genome from task B develops a state-1 generator proportional to the increase in the initial Life-pattern density, but it saturates in higher densities. Also, the evolved genome from task C has the capability to suppress the density of the initial Life pattern if it is high. However, such suppression is competing with the original Game of Life and thus is only effective for lower initial density cases.

Concerning the histograms in Fig. 7, we notice that there are two attractors induced by the second layer: one associated with very low-density states and the other associated with states around a density of 0.15. In fig. 8, there are also two attractors created by the genome activity: one near to a density value of 0 and the other near to a density value around 0.1.

How are these attractors chosen from the initial densities? Figure 10 and 11 are snapshots of genomes evolved from tasks B and C in the two initial density conditions. They reveal the appearance of the two attractors due to a subtle difference found in the starting configuration. For example, in Fig. 10, the two snapshots at t=1 have only a few bits which differ from each other, yet the higher-density initial pattern only makes state-1 clouds, thus creating a high-density resultant pattern.

This occurs because in the task B genome, the state-1 generator is only activated by few state-1 cells initially located in the space. Without these cells, the genome does not activate this generator and almost all the cells remain at state-0. Higher initial densities increase the probability that these cells become state-1. So the resultant densities are proportional to the initial density.

Similarly, the task C genome has a state-1 generator which is only activated when initial density is low. However, when only a few specific cells are state-1, the generator stops creating cloud-like patterns. Higher initial densities increase the probability of this event. Consequently, the resultant densities are inversely proportional to the initial density. When the initial density is higher than 0.1, the genome makes another spontaneous generator derived from the Game of Life, resulting in higher densities in a similar

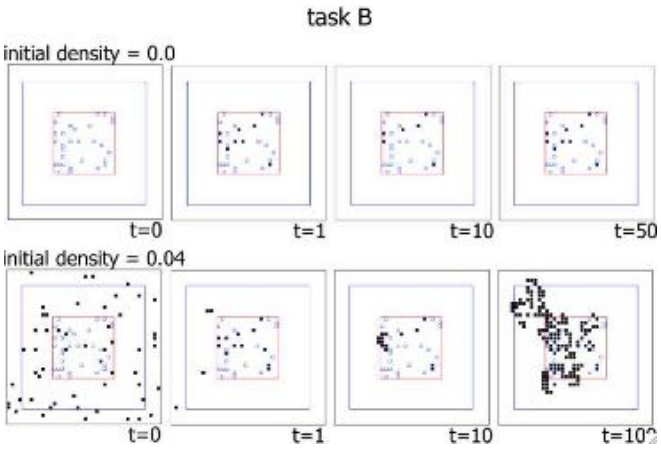

Figure 10: Snapshots of the Life dynamics of the fittest genome in task B. The top columns show the results when the initial density is zero, whereas in the bottom column the initial density is 0.04. In the zero-density condition, the Life pattern stays almost at zero activity. However, in the higher-density condition, a clouded pattern of state-1 cells emerges after t=10.

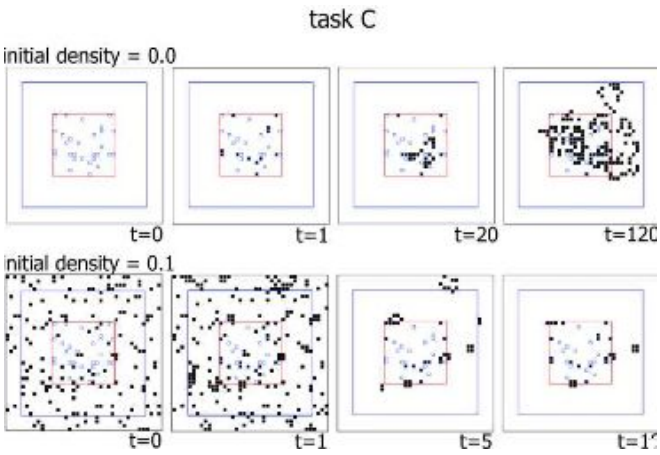

Figure 11: Snapshots of the Life dynamics of the fittest genome in task C. The top columns show the results when the initial density is zero, whereas in the bottom column the initial density is 0.1. In the zero-density condition, a clouded Life pattern emerges, but in the 0.1 condition, the initial state-1 cells eventually disappear.

manner to task B.

Discussion

In this paper, we studied homeostasis and adaptability with respect to cell states in the Game of Life as opposed to temperature as in Daisy world. There are some lessons to be learned here, both about homeostasis and the effect of noise.

The CA rule evolved in task A has the capability to generate a state-1 density regardless of the initial Life density. The behavior of this rule can be compared to a "random generator" which randomly updates the cell states at a certain probability. In task A, the evolved CA ruleset mimics a random generator using a deterministic rule. In tasks B and C there also have evolved generators of state-1 cells, but they are more sensitive to the initial state density.

Since we do not have any external noise in this simulation, it can be compared to deterministic chaos in continuous-state dynamical systems. Comparing the upper and lower figures in Fig. 10, we notice that almost identical initial Life patterns lead to different attractors. Such sensitivity to minute differences is also reminiscent of chaotic dynamics. Now we will discuss the nature of those two attractors in greater detail.

The initial Life states are evolved into either lower-density or higher-density states. Those two types of attractors are controlled by the evolved CA rules. We assume that the evolved CA ruleset may have both activators and inhibitors. Some CA rules tend to generate more state-1 Life patterns than the others and increase the state-1 density as the result, which we call "activator" rulesets. In contrast, some CA rules show the opposite behavior and lower the state-1 density, which we call "inhibitor" rulesets.

Daisy World	**Our model**
temperature	state-1 density
black daisy	activator CA
white daisy	inhibitor CA
growth rate as a function of the temperature	generators as a function of LIFE pattern

Table 2: A comparative chart between Daisy world model and the present LIFE game model

These two opposite behaviors remind us of the black and white daisies in Daisy World. Both these rulesets and those daisies can cooperatively make homeostatic states. Because black and white daisies have opposite responses to the sunlight, they can self-regulate the temperature by tuning their population size. If there are more black daisies, the temperature goes up as the average albedo gets lower, whereas if there are more white daisies, the temperature goes down due to the higher albedo value.

This simple scenario is also realized in the present Life game system. We simply take the activator rulesets as black daisies and the inhibitor rulesets as white daisies. The correspondence between Daisy world and this Game of Life system is shown in Table 2.

In Daisy world, the growth rates of daisies are determined by the local temperatures. The concept of temperature is not implemented in our system explicitly. Instead, local Life patterns determine the behavior of the CA rule sets. Note that the equivalent of Daisy world's local temperature in

our system is not just a one-dimensional variable which explicitly specifies the growth rate of state-1, but is instead a spatio-temporal pattern which drives responses from the evolved CA rulesets. This dynamical property of the Daisy World has also been discussed in our previous model of the mobile daisy agent(Ikegami and Suzuki, 2008).

Ideally, task A should be achieved by coupling the evolved CA rules in task B and C, however, we have not completed that task in this paper. Detailed analysis of that work will be reported in a follow-up to this research.

Instead of using chaos-like attractor in this study, it may be interesting to use unique Game of Life creatures such as oscillators, breeders and guns to generate homeostasis in more complex ways. We might then expect to see alternative homeodynamic mechanisms which are very different from those seen in Daisy world.

Acknowledgments

We thank Dr. Eric Silverman for his intense proof-reading of the manuscript. This work is partially supported by Grant-in aid (No. 19300104)fromtheMinistryofEducation, Science, Sports andCulture

References

Ashby, W. (1960). *Design for a Brain: The Origin of Adaptive Behaviour*. Chapman and Hall, London.

Cannon, W. B. (1932). *The wisdom of the body*. W.W. Norton and Company, New York.

Harvey, I. (2004). Homeostasis and rein control: From daisyworld to active perception. In Pollack, J., Bedau, M., Husbands, P., Ikegami, T., and Watson, R. A., editors, *Artificial Life IX : Proceeding of the 9th International Conference on the Simulation and Synthesis of Living Systems*, pages 309–314. Cambridge, MA: MIT Press.

Ikegami, T. and Suzuki, K. (2008). From homeostatic to homeo dynamic self. *BioSystems*, 91(2):388–400.

Lovelock, J. E. (1972). Gaia as seen through the atmosphere. *Atmos. Environ.*, 6:579–580.

Taylor, T. (2004). Redrawing the boundary between organism and environment. In Pollack, J., Bedau, M., Husbands, P., Ikegami, T., and Watson, R. A., editors, *Artificial Life IX : Proceeding of the 9th International Conference on the Simulation and Synthesis of Living Systems*, pages 268–273. Cambridge, MA: MIT Press.

Watson, A. J. and Lovelock, J. E. (1983). Biological homeostasis of the global environment: the parable of daisyworld. *Tellus*, 35B:284–289.

Wiener, N. (1948). *Cybernetics or Control and Communication in the Animal and the Machine*. Paris: Hermann et Cie - MIT Press, Cambridge, MA.