Artificial Life as Controlled Disequilibrium

Steve Battle
Computer Science and Creative Technologies
University of the West of England, Bristol
Bristol, UK
steve.battle@uwe.ac.uk

***Abstract*—Artificial Life can be modeled using simulated autopoiesis. Liquid Automata are used to define simulated chemical reaction systems; particle systems with rules governing how particles are transformed on collision with each other. Unlike cellular automata, there is no fixed grid or time-step, only particles moving about and colliding in continuous space/time. These systems may give rise to emergent artificial life, or they may be (artificially) lifeless. Can we distinguish between these systems by analysing their equilibria?**

I. Introduction

The dynamics of non-living systems is characterised by increasing entropy. A stone rolling to the bottom of a valley can be explained by purely physical forces that maximise entropy. By comparison, a life-form like a single-celled bacterium exploits external energy sources to minimise entropy locally, creating "order from disorder" far from chemical equilibrium [1]. According to James Lovelock, this principle can be used to look for life on alien worlds by analysis of planetary atmospheres [2]. A system that is far from equilibrium may provide evidence of life. On Mars, Curiosity rover discovered levels of oxygen 30% higher in spring and summer than expected. This may be evidence of life, or must be explained by other means.

II. Artificial Life

How then are we to understand life? Is it just a smorgasboard of features such as the capacity for growth, reproduction, functional activity, and change? This approach is not particularly scientific. Maturana and Varela argue that life is no more and no less than a system that exhibits *autopoiesis* [3], the ability of a living system to self-organise and produce itself in the physical realm. A living system is a self-referential domain of interactions in the physical space, generally a network of 'chemical' relationships. According to Maturana and Varela, "An autopoietic machine is a machine organized (defined as a unity) as a network of processes of production (transformation and destruction) of components." Yet, there are many kinds of chemical networks that aren't alive, consider a chemical explosion characterised by a runaway chain reaction of positive feedback. The signature of life is the emergence of structure that distinguishes self from non-self, closing it off from its environment, "A universe comes into being when a space is severed into two. A unity is defined." This closure emerges from, and is dynamically maintained by the organism. This principle is also central to Spencer-Brown's 'Laws of Form',

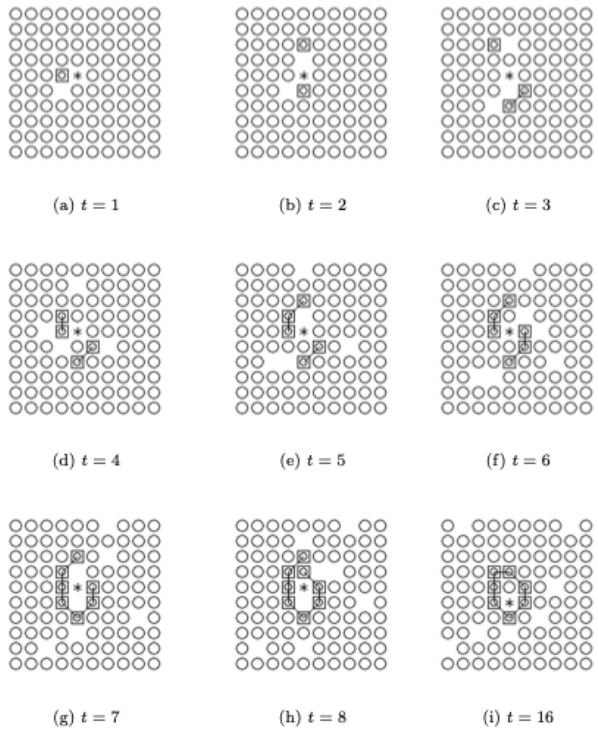

Fig. 1: A simulation of autopoiesis using a discrete time cellular automaton on a rectangular grid, based on Varela's original algorithm. In (a) a pair of substrate (circles) are transformed into a single link (squared circle) by the catalyst (asterisk). By (c) we see the first bonds forming, then in (i) these finally form a closed boundary around the catalyst.

"a distinction is drawn by arranging a boundary with separate sides so that a point on one side cannot reach the other side without crossing the boundary. For example, in a plane space a circle draws a distinction." [4]

There are two key features of autopoiesis:

1) Organisational closure: A network of the processes of production, e.g. a chemical reaction system.
2) Structural closure: The appearance of a homeostatically maintained boundary that divides self from non-self, also described as the maintenance of identity, e.g. a cell-wall.

We can simulate autopoietic processes computationally — not life as we know it, but *artificial* life. Conway's "Game of

Life" [5], [6] is a classic Cellular Automaton that exhibits emergent, self-organised behaviour. It gives rise to cyclic patterns that are self-reproducing, such as the 'blinker' and 'glider' patterns. But these patterns are not obviously life-like in the sense defined by Maturana and Varela as they lack a clear structural boundary (however, see Beer [7] for an alternative view). To demonstrate the process of autopoiesis, Francisco Varela devised a novel Cellular Automaton that shows how a simple system of rules can give rise to a 'cellular' structure [8]. It implements a simplified model of the chemical reaction rules found in living cells, demonstrating how these work together to form an organisationally closed system of production. Figure 1 shows output from this system; a catalyst (shown as an asterisk) is the 'seed'; substrate particles (circles) are consumed by the system, as pairs of them are composed by the catalyst to form new link particles (squared circles); and link particles self assemble into a compound structure — a primitive cell wall with bonds shown as lines drawn between them.

The rules implemented by Varela's Cellular Automaton are described as chemical reaction rules. A reaction rule has a left-hand side defining the *reactants*, separated by an arrow from the reaction *products* on the right-hand side. The appearance of a + (plus) symbol between reactants indicates an event where all of the indicated particles must come together. Each particle type may be prefixed by a positive integer indicating a number of particles of the same type, so that $K + 2S$ is equivalent to $K + S + S$, the interaction of three particles. The use of plus between reaction products, indicates that the reaction produces multiple outputs. Bonds are indicated by multiplicative operators such as $L.L$, L^2, or L^n [9]. These reaction rules are summarised below, where K represents the catalyst, S the substrate, and L the link particles.

$$\text{composition} : K + 2S \rightarrow K + L \quad (1)$$

$$\text{disintegration} : L \rightarrow 2S \quad (2)$$

$$\text{concatenation} : L^n + L \rightarrow L^{n+1} \quad (3)$$

Figure 1 illustrates a number of steps from a single run. In step (a) at time, $t = 1$, a pair of substrate particles are composed into a single link particle by the catalyst. After the composition of a number of link particles, we see the first bonds forming between them in step (c) at $t = 3$. The concatenation rules in Varela's algorithm are constrained to forming only obtuse bond angles. This prevents the uppermost link particle at (h) $t = 8$, from bonding with the particle immediately below it. It is only when this particle disintegrates later at (i) $t = 16$, that enables the remaining links to re-bond, and form a closed boundary around the catalyst.

These same rules were re-implemented in a later program called SCL (Substrate, Catalyst and Link) by McMullin using the SWARM system [10] where chemical reaction rules are captured in a modular fashion enabling their reaction rates to be more precisely controlled. This program introduced a more configurable way to control the random motion of particles in 2D space, modelled as a square lattice with toroidal toplogy and discrete time.

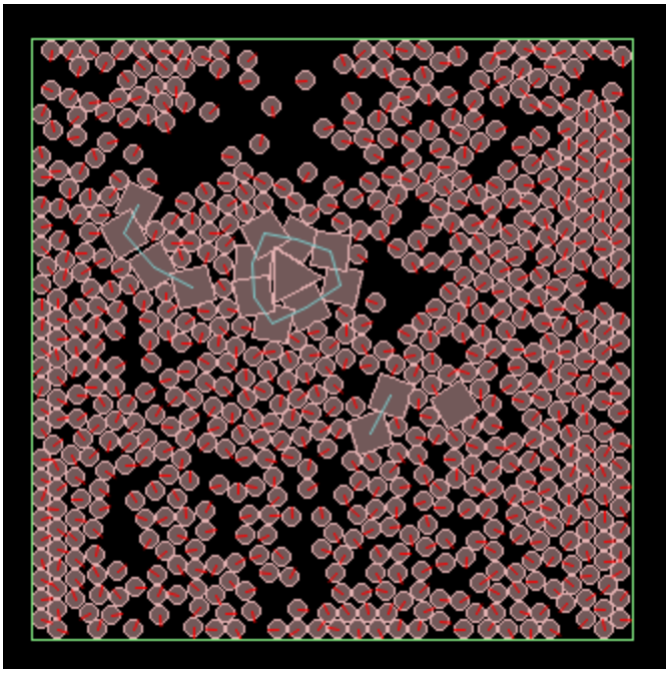

Fig. 2: Liquid automaton showing a boundary (blue links) forming around the catalyst (triangle), distinguishing self from non-self. The catalyst transforms the substrate (circles) into its structural building blocks (squares).

III. Particle Systems

There is currently a great deal of interest in 2D particle systems, though the same ideas extend to three dimensions. Particle systems are game physics engines designed to reproduce naturalistic phenomena based on objects moving around, typically in a 2D space. The engine used for the Liquid Automata described in this paper is Box2D (specifically pybox2D), a rigid body simulation library for games. Interestingly, Box2D has been used as the game engine for a number of implementations of "Angry Birds." Each particle is a 2D body with mass and velocity, so particles have three degrees of freedom; translation along x,y axes, and rotation. Each body is associated with one or more shapes which can be any geometrical construct, such as the squares, circles, and triangles seen in Figure 2. Joints define constraints on the relative motion between two bodies, used here to create bonds between neighbouring links. Forces, torques, and impulses are applied to bodies to make them move. Box2D includes a high performance iterative constraint solver that resolves joint constraints, particle motion and resulting collisions [11]. Particles can bounce off each other in elastic collisions, or slide against each other based on a realistic simulation of the frictional forces between them. A world may, or may not, have gravity; the simulations described here do not.

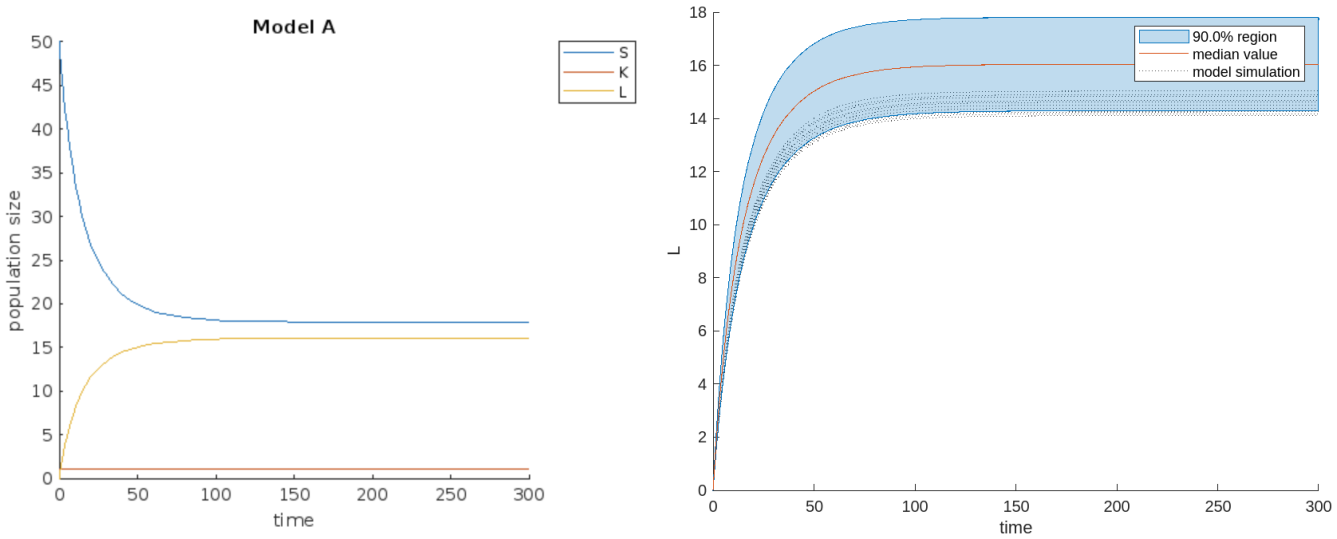

(a) (b)

Fig. 3: (a) Model A runs to equilibrium with L converging at around 16. (b) This equilibrium is sensitive to variations in initial population size, S. 'Sobol' analysis shows the effect of this variation, with the blue region covering 90% of the observed variation.

IV. Liquid Automata

We reproduce and extend Varela's simulations of autopoiesis [12] using so-called *Liquid Automata*. These are based on 2D particle systems with added support for chemical reaction rules. Unlike cellular automata which divide space into to a regular fixed grid, and time into discrete steps, Liquid Automata implement a continuous model of space/time. There is no fixed grid, just particles moving about and colliding with each other in continuous space. By analogy with cellular automata, the Liquid Automaton is a variety of *collision-based* system [13]. Particles move freely in space and react 'chemically' when they collide. Rules (1) & (3) are invoked when the necessary particles come into contact with each other, while rule (2) occurs spontaneously.

External energy is added to the system in the form of random 'Brownian' motion defined in terms of a Wiener process along the x,y dimensions [14]. The force applied to every particle along each dimension is a normally distributed random variable with zero mean, and variance, $(delta)^2 dt$, correlated with a single parameter delta, and time period, dt, which varies dynamically. It is this external energy source that drives the system to self-organise, and create a local island of order from the surrounding chaos.

A Liquid Automaton defines an organisationally closed chemical reaction system (CRS). Figure 2 shows output from a Liquid Automaton implementing reaction rules (1), (2), and (3). Bathed in a liquid substrate, the catalytic agent (triangle) composes substrate (circles) into its structural building blocks, or links (squares). Links are subject to decay, and may spontaneously disintegrate back into a pair of substrate particles. The link particles are able to self-organise (blue links) into a structure akin to a long-chain polymer. Links can make up to two connections, and a chain of concatenated links is able to wrap around and close in on itself. This emergent closed boundary, analogous to a cell wall, divides self from non-self; the signature of (artificial) life.

V. Equilibrium and Disequilibrium

If artificial life is to be recognised by its disequilibrium, we must first define what the equilibrium state would be in the absence of artificial life. If a key characteristic of life is the boundaries that it forms, then we need to study the system in the absence of those boundaries. Deleting the concatenation rule (3) that bonds particles together, disrupts the internal organisation of the system. The network of relations no longer performs a complete cycle that produces an enclosed structure so the boundary never forms. We will call the this simpler 'chemical' system, System A, while the full system with boundaries will be known as System B. For experimentation we need a measurable proxy for this emergent structure. By observation, the boundary appears to regulate the number of substrate particles reaching the catalyst, with the knock-on effect of throttling back the number of link particles produced by the reaction. We therefore use the number of link particles as a proxy for this structure and *hypothesise that the mean number of link particles is higher in System A than it is in System B*.

The Liquid Automaton for System B maintains a mean of 12.46 link particles, based on samples taken every second over a period of 30 minutes, with an initial substrate population of 500 substrate particles. Deleting the concatenation rule in System A, we see an increased mean of 17.56 link particles over the same time period, given the same initial conditions.

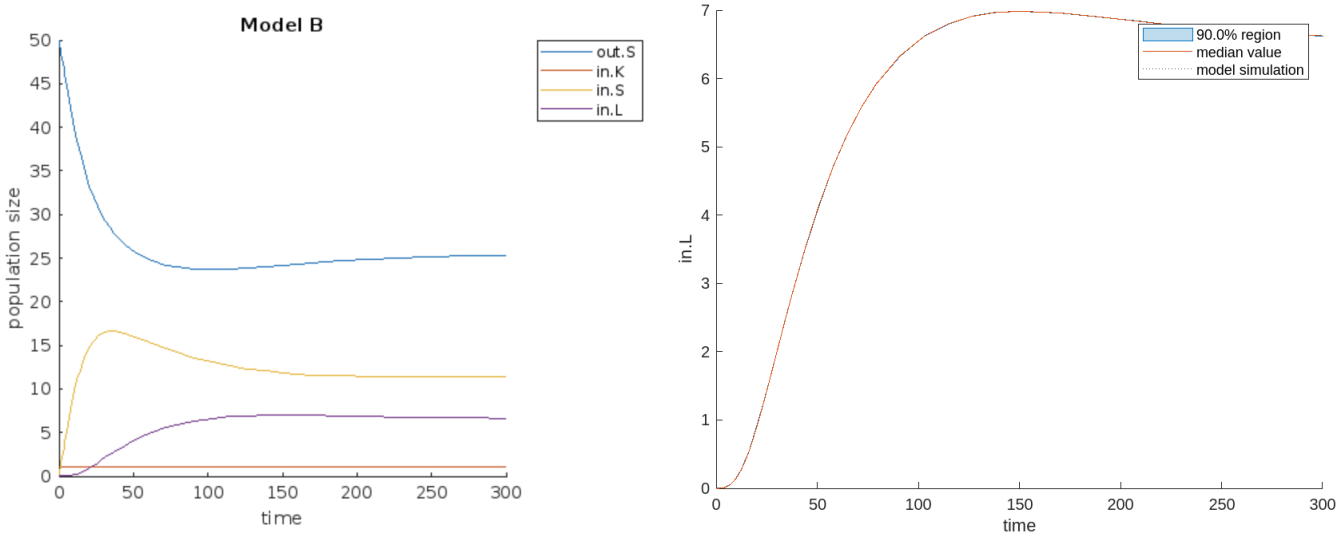

(a) (b)

Fig. 4: (a) Model B runs to equilibrium with $in.L$ converging around 6.6. There is an initial surge of substrate, $in.S$, into the interior before the boundary is fully formed. (b) The 'Sobol' analysis is absent of any variation, so the equilibrium of $in.L$ is independent of the initial population size, $out.S$.

This 29% difference from artificial chemistry is the 'tell' of artificial life at work. Its emergent structure shifts it away from 'chemical' equilibrium.

We further hypothesise that the higher point of equilibrium for System A is sensitive to initial conditions. Simulation with Liquid Automata is computationally expensive, and the data it produces is too variable to perform sensitivity analysis reliably. Even without the concatenation rule, the halo of reaction products around the catalyst produces a barrier that impedes subsequent reactions with the substrate. Instead, we define simplified models using the Mathworks SimBiology package [15]. These models are more tractable as they do not require a physics engine, and only take population size into account. They cannot produce any organised states of matter. We define a simpler Model A, corresponding to the Liquid Automaton of System A. Parameters governing the reaction rates are set to simple 'Mass Action' kinetics, where the rate of a chemical reaction is directly proportional to the concentration of the reactants. Model parameters governing reaction rates are tuned only to the extent of achieving comparable qualitative behaviour, with results of the same order of magnitude. Finer tuning doesn't appear to provide any additional insight.

Model A comprises just the reaction rules for composition (4) by the catalyst, K, and spontaneous disintegration (5) of link particles back into substrate, S, at the rates indicated. There is no rule for concatenation as the simulator does not support the notion of compound particles. We will continue to refer to 'link' particles for continuity with Varela's nomenclature, but these numerical simulations do not perform any such linking. This 'chemical' simulation will naturally run towards a state of equilibrium, at a rate governed by the "Law of Mass Action". These rules and their corresponding (mass-action) reaction rates are defined below.

$$\text{composition} : 2S + K \rightarrow K + L \qquad \text{rate = 0.5E-3} \qquad (4)$$

$$\text{disintegration} : L \rightarrow 2S \qquad \text{rate = 0.01} \qquad (5)$$

Figure 3(a) shows a simulation run of model A. In these simulations we only consider a single catalyst particle, K, remaining constant at 1. The simulations begin with no link particles, and it has an initial population size of $S = 50$. The system reaches equilibrium when the production and decay of link particles, L, reaches a balance, converging at around 16 particles. Sensitivity analysis allows us to explore the effects of variations in model quantities on a model response. We want to explore the time-dependent sensitivities of the number of link particles, L, with respect to the initial substrate population, S. Using Local Sensitivity Analysis (LSA) we compute, $\delta L/\delta S = 1.22$. This shows a clear sensitivity of the state of the output species L on the initial population size of the substrate, S. This is positive, showing that in Model A, L increases with S, as expected. Figure 3(b) illustrates the results of a variance-based sensitivity analysis using the 'Sobol' method, showing the simulation results and mean value. The shaded region shows the spread of outputs, covering 90% of the results, based on a random spread of perturbations to the initial population size. We conclude that the stability exhibited by Model A is contingent, being sensitive to initial conditions.

VI. Controlled Disequilibrium

By "controlled disequilibrium", we mean a new point of equilibrium displayed by a self-organising system that is both distinct from the 'chemical' equilibrium we find in the absence of self-organisation, and is intrinsically stable. This intrinsic stability can be understood as an internally set goal or purpose. In other words, it behaves as a control system to maintain that goal. As before, we measure the mean number of link particles observed from second to second and investigate whether the goal criteria, the goal value for L, is stable in the face of perturbations in the initial conditions. *We hypothesise that the artificially living system maintains this new point of equilibrium over a range of initial conditions*.

As before, we build a simplified model of System B, using MathWorks SimBiology [15], and call this Model B. Whereas in the System B Liquid Automaton, the boundary is an emergent property, Model B introduces explicit compartments that separate particles inside the boundary from those outside. The catalyst, K, now only reacts with substrate, S, inside the boundary (rule 6). We also assume that when link particles disintegrate, the waste substrate is 'dumped' outside the boundary, ready for recycling (rule 7). The boundary is semi-permeable allowing substrate particles to diffuse across in either direction at a rate that is inversely proportional to the density of link particles forming the boundary (rule 8). Model B incorporates the reaction rules below.

$$\text{composition} : 2\,in.S + in.K \rightarrow in.K + in.L \text{ rate=0.5E-3} \quad (6)$$

$$\text{disintegration} : in.L \rightarrow 2\,out.S \quad \text{rate=0.01} \quad (7)$$

$$\text{diffusion} : out.S \longleftrightarrow in.S \quad \text{rate}=1/(in.L+1) \quad (8)$$

Figure 4(a) shows a simulation run of model B, with an initial population size of $out.S = 50$. There is an initial surge of substrate into the interior before the boundary is fully formed (the density of link particles increases). The quantity of catalyst, $in.K$, again remains constant at 1. The mean number of link particles, $in.L$, converges at around 6.6 particles, lower than the point of equilibrium for Model A. Again, using Local Sensitivity Analysis (LSA), we compute, $\delta in.L/\delta out.S = 0$, the time-dependent sensitivities of the output species $in.L$ with respect to the initial substrate population, $out.S$. This is 0, demonstrating that $in.L$ does not vary with $out.S$, and so there is no sensitivity to perturbations in $out.S$ over the range tested. Similarly, the 'Sobol' plot in Figure 4(b) shows no variance in $in.L$. The shaded region of variance is all but invisible.

We are in effect applying Ashby's "test for independence" [16, p158] to demonstrate that Model B induces temporary independence (over the range tested) between the initial population of substrate, and the density of link particles; "constancy in a subsystem's state implies that the state is one of equilibrium, and constancy in the presence of small disturbances implies stability." We conclude that the stability exhibited by Model B is largely independent of initial conditions. The variable diffusion rate creates a feedback control loop based on the quantity of link particles, making it robust to disturbance.

VII. Conclusion

A central feature of living systems is that they exist out of equilibrium with their environment. All closed systems will ultimately reach thermodynamic or chemical equilibrium, but this may simply be contingent on initial conditions. A key aspect of this disequilibrium is that it is under control by the organism and maintained within a specific range necessary for the organism's continued survival.

Liquid Automata enable us to investigate the mechanisms of artificial life based on particle simulations with added reaction rules. They are a tool for exploring emergent phenomena and the equilibrium states that arise. The equilibrium state of artificial life (simulated autopoiesis) falls far from the baseline equilibrium of an (artificially) lifeless chemical reaction system. This baseline equilibrium is also contingent on initial conditions. Artificial life self-organises, defining a new point of disequilibrium relative to this baseline. Furthermore, this new equilibrium is dynamically controlled through feedback, remaining stable in the face of perturbation. Artificial Life has a goal, its purpose to survive.

References

[1] E. Schrodinger, *What is Life?* Cambridge, UK: Cambridge University Press, 1944.

[2] J. E. Lovelock, "A physical basis for life detection experiments," *Nature*, vol. 207, pp. 568–570, 08 1965.

[3] H. Maturana and F. Varela, *Autopoiesis and Cognition: The Realization of the Living*, ser. Boston Studies in the Philosophy and History of Science. Dordrecht, Holland: D. Reidel Publishing Company, 1980.

[4] G. Spencer-Brown, *Laws of Form*. London: George Allen and Unwin Ltd., 1969.

[5] M. Gardner, "The fantastic combinations of john conway's new solitaire game "life"," *Scientific American*, vol. 223, pp. 120–123, 1970.

[6] E. R. Berlekamp, J. H. Conway, and R. K. Guy, "Winning ways for your mathematical plays," *Journal of Multivariate Analysis*, 1982.

[7] R. D. Beer, "Characterizing autopoiesis in the game of life," *Artificial Life*, vol. 21, no. 1, pp. 1–19, 2015.

[8] F. Varela, *Principles of Biological Autonomy*. 52 Vanderbilt Avenue, New York, New York 10017: Elsevier North Holland, Inc., 1979.

[9] W. Hordijk and M. Steel, "Autocatalytic sets and boundaries," *Journal of Systems Chemistry*, vol. 6, no. 1, 2015.

[10] B. McMullin, "SCL: An artificial chemistry in Swarm," *Santa Fe Institute*, 02 1997.

[11] E. Catto, "Iterative dynamics with temporal coherence," *Game Developer Conference*, 03 2005.

[12] F. Varela, H. Maturana, and R. Uribe, "Autopoiesis: the organization of living systems, its characterization and a model," *Currents in modern biology*, vol. 5, no. 4, pp. 187–196, may 1974.

[13] T. Toffoli, "Symbol super colliders," in *Collision-Based Computing*, A. Adamatzky, Ed. London: Springer London, 2002, pp. 1–23.

[14] E. Nelson, *Dynamical theories of Brownian motion*. Princeton, New Jersey: Princeton university press, 1967.

[15] The MathWorks Inc., "Simbiology (r2020a)," Natick, Massachusetts, United States, 2020. [Online]. Available: https://uk.mathworks.com/help/simbio/

[16] W. Ashby, *Design for a brain*, ser. Science Paperbacks. London: John Wiley & Sons Inc., 1960.