Autopoiesis in RealLife Euclidean Automata

Thomas M. Gaul[1]
[1]Cognitive Science Program, Indiana University Bloomington
tgaul@iu.edu

Abstract

Autopoiesis aims to describe the organization and limits of living systems. Unfortunately, its theoretical development has largely been carried out verbally, with less focus on developing formal concepts of the key ideas, such as structure, organization, process, etc. Using toy models of emergent individuals allows us to fully characterize these concepts concretely. This paper generalizes previous work that analysed autopoiesis in the Game of Life. I use the Larger than Life family of cellular automata to explore how the concepts of production process, autopoietic network, and cognitive domain extend to this space, before moving to the continuum limit in RealLife — a continuous-space, discrete-time family of Euclidean automata.

Introduction

Molecular biology seeks to describe the molecules that compose living systems. Evolutionary theory seeks to describe the processes involved in the historical transformation of living systems. Cognitive science seeks to describe the mechanisms that underlay the behaviour of living systems. But what is a "living system?" — What is missing in these established fields is an account of the organization of living systems and how that organization is actively realized and maintained. That is, we want to understand how *emergent individuals* maintain their identity as distinct unities in space through their underlying constituent processes. Without such an account, we are left to assume the existence of living systems with little understanding of what it actually means to do so.

Maturana and Varela's notion of *autopoiesis* provides a compelling way to think about the organization of living systems (Maturana and Varela, 1980). An autopoietic system is a network of processes of production that produce components that both realize that same network (self-production) and constitute it as a distinct unity in space (self-distinction). Despite its relatively succinct definition, the idea is quite powerful in its consequences for understanding cognition, evolution, and epistemology (Varela, 1979; Maturana and Varela, 1980, 1987) — however, there is still much ambiguity and controversy in the details (e.g., Di Paolo, 2005; Virgo et al., 2011).

To further the theoretical development of autopoiesis, then, it is useful to use toy models of emergent individuals as a way to work out these ideas in a concrete instance. Such models have been around for decades now (Varela et al., 1974; McMullin, 2004), but they have been used more as demonstrations than as opportunities for deep analysis. More recently, however, Beer has used gliders in the Game of Life (GoL) cellular automaton (Adamatzky, 2010) to thoroughly investigate autopoiesis and its consequences (Beer, 2014, 2015, 2020a,b,c). This paper presents an initial foray extending this work into continuous space. Specifically, I use a generalization of Game of Life, called Larger than Life (Evans, 2001), and its continuum limit, RealLife (Pivato, 2007). This provides a more incremental path to fully continuous universes, such as Lenia (Chan, 2019), since RealLife is still discrete in time while being continuous in space.

Thus, this paper proceeds in three stages: first, an introduction to the Larger than Life (LtL) family of cellular automata; second, how the notions of autopoiesis and cognition, as articulated in the GoL literature, extend to LtL and the problems that arise with that extension; and third, how these problems scale when taken to the continuum limit in RealLife. To alleviate some of these problems, I propose an object — the *density manifold* — as a way to formulate and visualize constraints on the organization of emergent individuals in order to simplify the analysis of structures and their interactions.

Larger Than Life

LtL is a family of outer totalistic cellular automata (CA) defined on a two-dimensional lattice $\mathbb{Z}^2$ where the state of the LtL universe is given by a function $\mathbf{u} : \mathbb{Z}^2 \to \mathcal{M}$, where $\mathcal{M} := \{0, 1\}$. Each lattice cell has a neighbourhood; here, we will use the generalized Moore neighbourhood defined $\mathcal{N}_\rho^x := \{x - y \,:\, ||y||_\infty \leq \rho, y \in \mathbb{Z}^2\}$, where $\rho \in \mathbb{N}$ is the neighbourhood radius, $x \in \mathbb{Z}^2$ is the location of the center cell, and $||\cdot||_\infty$ is the max-norm. We can define a convolution kernel κ to give the density of on-cells in the neighbourhood of a cell at some point on the lattice x:

$$\kappa(y) := |\mathcal{N}_\rho|^{-1}\mathbb{1}_{\mathcal{N}_\rho}(y) \tag{1}$$

$$\kappa * \mathbf{u}(x) = |\mathcal{N}_\rho|^{-1}\sum_{y\in\mathcal{N}_\rho}\mathbf{u}(x-y). \tag{2}$$

Then a rule $(\rho, b_0, b_1, s_0, s_1)$, with birth interval $[b_0, b_1]$ and survival interval $[s_0, s_1]$ such that

$$0 \le s_0 \le b_0 \le b_1 \le s_1 \le 1,$$

defines an LtL CA $\xi : \mathcal{M}^{\mathbb{Z}^2} \to \mathcal{M}^{\mathbb{Z}^2}$:

$$\begin{aligned}\xi(\mathbf{u})(x) :=\, &\mathbf{u}(x)\cdot\mathbb{1}_{[s_0,s_1]}(\kappa*\mathbf{u}(x))\\ &+(1-\mathbf{u}(x))\cdot\mathbb{1}_{[b_0,b_1]}(\kappa*\mathbf{u}(x))\end{aligned} \tag{3}$$

where $\mathbb{1}_{\mathscr{A}}$ is the indicator function of a set $\mathscr{A}$. Thus, GoL is the LtL rule $\left(1, \frac{3}{9}, \frac{3}{9}, \frac{3}{9}, \frac{4}{9}\right)$ that gives rise to what can be called the *Conway physics*. If we interpret the elements of $\mathcal{M}$ as *components*, then the Conway physics in turn gives rise to a spatial chemistry on 0- and 1-components (off- and on- cells) defined by a rule for transforming reactants to products: $\mathbf{u}|_{\mathcal{N}_\rho^x} \mapsto \xi(\mathbf{u})(x)$ (Beer, 2015).

While GoL is a single rule with interesting dynamics that support many recurrent spatio-temporal patterns (Adamatzky, 2010), it is a non-trivial task to find other LtL rules with similar properties (Evans, 2001). Moreover, scaling such rules to arbitrary neighbourhood size does not guarantee similar dynamics at those sizes. There is also the problem of determining the existence of stable structures for arbitrary rules (Evans, 2003).

There are three general classes of such structures that we are interested in here: *still-lifes*, *oscillators*, and *bugs*, generalizing the *block*, *blinker*, and *glider*, respectively, from GoL. *Still-lifes* are fixed points of the dynamics, *oscillators* are limit-cycles, and *bugs* are periodic mod-translation. Evans (2003) has provided rigorous results for the existence of still-lifes and oscillators, as well as empirical results that suggest that bugs can also be scaled arbitrarily. I will here focus on the block generalizations, as they are the simplest structures to analyse and are the only finite ones that have been formally proven to exist in the continuum limit (Pivato, 2007, pp. 62).

An LtL rule $(\rho, b_0, b_1, s_0, s_1)$ supports a still-life that generalizes the block when $s_0 \le (\rho+1)^2/|\mathcal{N}_\rho| \le s_1$ and either:

$$\text{(i)}\quad b_0 > \frac{\rho(\rho+1)}{|\mathcal{N}_\rho|}; \qquad \text{or}$$

$$\text{(ii)}\quad b_0 \le \frac{n}{|\mathcal{N}_\rho|} \le b_1$$

such that $n < \rho(\rho+1)$ and $n \ne ab$ where $a, b \in [0, \rho+1] \subset \mathbb{Z}$ (Evans, 2003, pp. 64). Any such rule would be sufficient for purposes of demonstration, but I will use *Bosco*'s rule $\left(5, \frac{34}{121}, \frac{45}{121}, \frac{33}{121}, \frac{57}{121}\right)$ as its dynamics are chaotic, similar to GoL, and it supports many analogous structures as well, including a 6×6 block (Figure 1; Evans, 2001). In general, though, the following will apply to any $(\rho+1)\times(\rho+1)$ block.

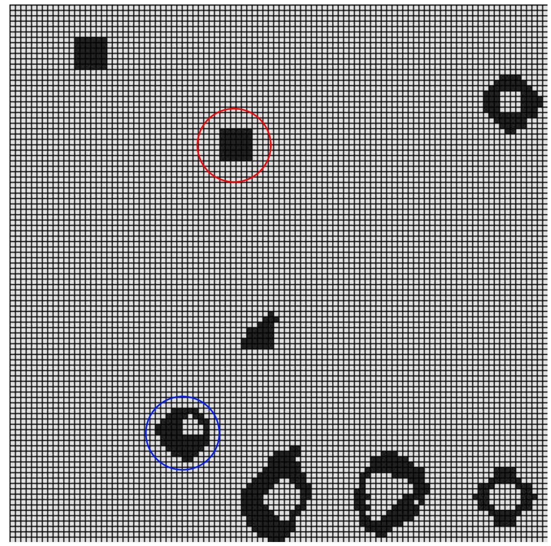

Figure 1: Example of entities supported by *Bosco*'s rule in a 100×100 periodic universe. In red is circled a 6×6 block. In blue is circled a bug. Black cells are on, white cells are off.

Autopoiesis

Now to proceed to the first step of the analysis: what components, exactly, constitute the block? We must first recognize that a block cannot be identified *de novo* in arbitrary universes. This is because we are considering the block as a *recurrent* pattern of production processes. If all the blocks we observe are immediately destroyed after they are realized, we have no way of recognizing their recurrent (closed) nature. However, once we develop a criterion for distinguishing blocks, it then becomes possible to identify them in arbitrary universes irrespective of whether their closure is realized.

In what universes, then, can we distinguish a block? Certainly, we will consider the 1-components usually associated with any recurrent pattern as belonging to the block; I will denote this set $\mathcal{B}^0$. From this, we then want to identify those components that support the closure of this pattern. Following Beer (2004), I identify these components as belonging to the union of the neighbourhoods of all the 1-components; that is, all the components that can immediately influence $\mathcal{B}^0$. The ρ-wide buffer of 0-components between $\mathcal{B}^0$ and the environment then constitutes the boundary of the block (Figure 2); let this region be denoted $\mathcal{B}^1$. This boundary acts as the interface between the environment and the block's internal components.

The identification of a block as $\mathcal{B}^{(0,1)}$ splits the LtL uni-

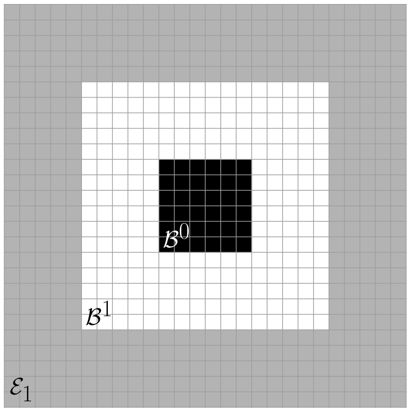

Figure 2: 6×6 block supported by *Bosco*'s rule. Black cells are 1-components in the region denoted $\mathcal{B}^0$; white cells are 0-components in the region denoted $\mathcal{B}^1$; and grey cells are unspecified environmental components in the region denoted $\mathcal{E}_1$.

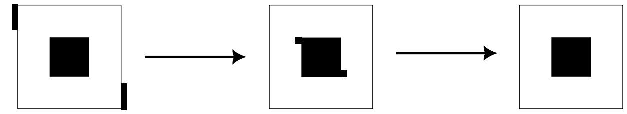

Figure 3: Example of structural degeneracy in the 6×6 block. The black outline demarcates the components belonging to the block. The arrows indicate updates to the LtL universe.

verse in two: block and not-block. But we can further decompose the not-block region based on the speed at which information can travel from the environment to $\mathcal{B}^{(0,1)}$. Since the maximum rate of information travel is ρ per unit time, the immediate surrounding ring of width ρ constitutes the complete set of components that can influence the block's state at the next step. The ring immediately surrounding *this* is the complete set of components that can only influence the block in at least two steps in time. Beer (2014) calls these rings the 1-environment and 2-environment, respectively. I will notate these as $\mathcal{E}_i$ for $i \in \mathbb{N}$.

Our decision as to what counts as a block defines the conditions under which the block exists — what configurations of components are *viable*. If we consider $\mathbb{1}_{\mathcal{B}^0}$ to be the only viable form of the block, then there is a complete lack of structural degeneracy: the structure cannot change without the block disintegrating. However, we may also wish to consider, for example, a block with two 1-components in the boundary as viable, since this structure returns to the canonical instance in one step (Figure 3). This is akin to the 'rocket' and 'wedge' forms of the glider in GoL, in that the two structures are both viable and can be transformed into one another (Beer, 2004). Note also that as the number of block components increases with ρ, so does the number of possible structures. Then, with the inclusion of this degeneracy, our definition of the block turns into a *set* of structures, each with its own constraints delineating what perturbations will destroy it or transform it into another viable structure (possibly itself). This set is partially observer dependent, though, as deciding how many steps constitute an acceptable "recovery time" is, to a certain extent, arbitrary.

To characterize a block in autopoietic terms, we need to define what a production process is and how these processes depend on one another. Here, the chemical metaphor is useful to relate local events in LtL to the metabolism and self-construction of cells that is usually thought of as the canonical instance of autopoiesis (Beer, 2015). This leads us to conceive of processes as relations between reactant and product components mediated by the update rule.

Due to the discrete-time nature of LtL, it is rather straightforward to carry over Beer's definition to the more general case. A process then becomes a $(2\rho + 1) \times (2\rho + 1)$ spatial configuration of components. The product $\xi(\mathbf{u})(x)$ becomes a function of the central component and the density of 1-components in its neighbourhood (its '1-density'). We can then sort these processes into four classes (Beer, 2015, pp. 5): if the central component is 0, the process is a *production* when $\kappa * \mathbf{u} \in [b_0, b_1]$ and is a *0-maintenance* otherwise; if the central component is 1, the process is a *1-maintenance* when $\kappa * \mathbf{u} \in [s_0, s_1]$ and is *destruction* otherwise.

We can further define a *partial process* (Figure 4) as an equivalence class of processes that agree on a subset of components, implying that the realization of such a process is contingent on the particular environment in which it is embedded (Beer, 2020b, pp. 204). These occur in the block when $\mathcal{N}_\rho^x$ intersects the environment. Note that the product of a partial process may be fully determined by only a subset of the neighbourhood when the birth and survival intervals contain more than one element. Another interesting possibility — dependent on structural degeneracy — is the occurrence of partial destruction processes. For example, the 1-components in the boundary in Figure 3 are normally destruction processes, but may be sufficiently perturbed by the environment and thus become 1-maintenance processes.

The dependency between two processes is a temporal relation realized via the components they produce (Beer, 2015). Since a process is instantiated by a particular spatial configuration of components, it is dependent upon another when the product of a preceding process is part of that spatial configuration (Figure 5). As neighbourhood size increases, dependency relations give greater spatial detail. In addition, the neighbourhood constraints implied by a given process become even clearer. Figure 6 shows an example. Two adjacent processes in the block always share *at least* $|\mathcal{N}_\rho| - (4\rho + 1)$ of their components, and as ρ approaches ∞, the proportion of components shared between those pro-

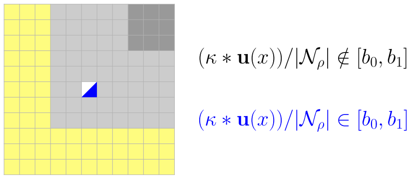

Figure 4: A partial process with $\rho = 5$. Yellow cells are unspecified, light grey cells are 0-components, and dark grey cells are 1-components. The half-white, half-blue square indicates that this process could be 0-maintenance or production. The corresponding equations specify the conditions under which the process is of either kind.

cesses approaches 1.

The point of defining processes and dependency as I have is to construct the *dependency graph*, representing the autopoietic organization as a closed network of production processes (Beer, 2015, 2020b; Beer et al., 2024). Extracting this representation is fairly simple: enumerate all the dependencies between processes in sequential block configurations. In the case of the block, dependency relations are equivalent to the spatial relations between processes, since no translation occurs. An example of dependency links in LtL is given in Figure 5b.

The Cognitive Domain

From this organization, we can derive the *interaction graph* of the block, describing perturbations in $\mathcal{E}_1$ and the transitions they induce (Beer, 2014; Beer et al., 2024). These perturbations are defined in terms of the block's response to them, and as such the classes of perturbations differentiated by the block constitute an *Umwelt* (von Uexküll, 1992). Maturana and Varela call this space of non-destructive perturbations the *cognitive domain* (Maturana and Varela, 1980, 1987).

The derivation is possible due to the neighbourhood constraints implied by the spatial configurations defining any given process, thus permitting a spatial reconstruction of all possible process sets that result in valid blocks when embedded in space. The dependencies between sets of processes can then be collapsed into edges representing transitions between viable structures (Figure 7).

There are two ways to analyse the structure of perturbations: (i) explicitly enumerate all perturbations and the block's response to them (Beer, 2014); or (ii) derive constraints from the organization that define what conditions a perturbation must satisfy to induce a given transition (Beer, 2020b). This latter option is the more preferable, especially as both the space of viable structures and possible perturbations increase with ρ. Thus, we will want to define an object on which these constraints can be defined: the *density grid* (Figure 8). The derivation of the grid from the dependency graph is simple: replace the processes of a spatially em-

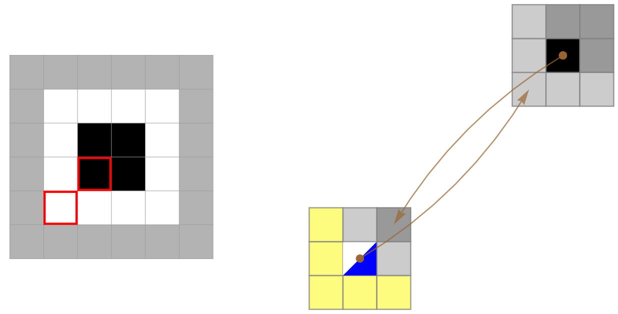

(a)

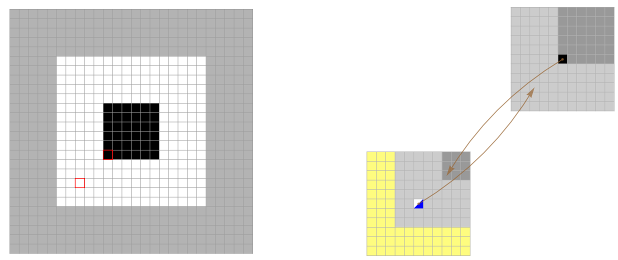

(b)

Figure 5: Process dependency in GoL and LtL. Red squares indicate where the processes are in the block. Black indicates a 1-maintenance process, white a 0-maintenance process, and blue a production process. Brown arrows indicate dependency relations such that the product of one process is a reactant in the other. Only two dependency links are shown, but there should be a link for every reactant in $\mathcal{B}^{(0,1)}$. (a) Dependency links in the 2×2 GoL block. (b) Dependency links in the 6×6 block supported by *Bosco*'s rule.

bedded graph with the 1-densities of those processes. The grid can also be generated directly from the universe by Equation (2). For compactness of notation, we will renotate Equation (2) as an instance of the more general density function Ψ.

The density grid can be used to derive perturbation classes by setting constraints on the 1-densities of each viable structure such that transitions to other viable structures are induced. For example, the constraint that induces a transition from the canonical block to itself is shown in Figure 7. Note that the constraint need only be specified on the boundary $\mathcal{B}^1$, since only those components are directly influenced by the environment. Thus, the density grid over the internal components is always invariant with respect to any particular viable structure.

Scaling

A few observations can be made about how autopoiesis generalizes to LtL. First, the significance of individual compo-

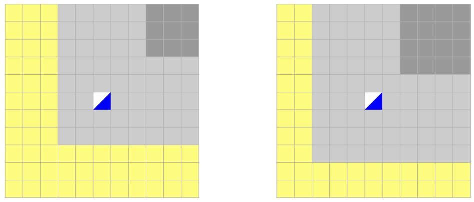

Figure 6: Two adjacent processes in the block. Yellow cells are unspecified, light grey cells are 0-components, and dark grey cells are 1-components. The process on the right is $(1, 1)$ from the process on the left.

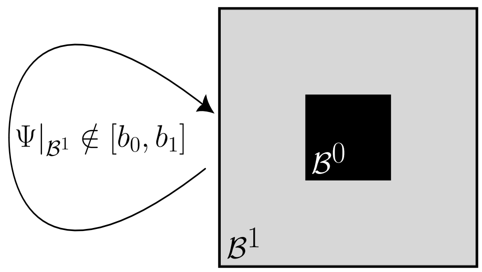

Figure 7: Interaction graph of a block without structural degeneracy. Ψ is the density function.

nents scales inversely with ρ (at least when the birth and survival intervals also increase with ρ). What this means is that it becomes more and more useful to consider components and processes in sets with measure corresponding to the area they cover. Second, increasing neighbourhood size raises concerns about the locality of the dynamics. We can resolve this point by noting that, if we conceive LtL universes as existing in $\mathbb{R}^2$, scaling the size of cells in proportion to $(\rho + 4)^{-2}$ ensures blocks of different radii can be contained within the same *area*. Thus, in a certain sense, we can consider the neighbourhood to be a constant region in $\mathbb{R}^2$ while we change the spatial density of cells in the universe.

Constructing the dependency graph is a matter of enumerating the dependency links between processes, and constructing the interaction graph is a matter of enumerating transitions between viable structures. Moreover, in both objects, increasing structural degeneracy adds more vertices to the graphs, corresponding to the number of viable structures or the number of processes composing those structures.

To estimate this scaling, we will approximate the number of viable structures by assuming that the preimages of a given configuration scale similarly with size as in GoL (Beer, 2017). Let $f(x^n)$ be a real polynomial function of x with degree at most n; then we estimate the number of preimages of a block by $2^{f(\rho^2)}$. It is more difficult to be specific about the number of edges in the interaction graph, but we can get a rough approximation by assuming that, on average, each configuration has a viable transition to some constant fraction γ of all other configurations. Thus, $\gamma 2^{f(\rho^2)} 2^{f(\rho^2)} = \gamma 2^{2f(\rho^2)}$ provides an estimate of the number of edges. If we assume all viable configurations have, on average, the same number of components, then the number of vertices in the dependency graph is roughly $(3\rho+1)^2 2^{f(\rho^2)}$. For the edges, we will assume that the average number of dependency links in a viable transition is the same as the number of dependency links between the canonical block and itself. This turns out to be $\frac{1}{2}(5\rho^2 + 6\rho + 1)$ when ρ is odd, making the total number of edges in the dependency graph approximately $\frac{\gamma}{2}(5\rho^2 + 6\rho + 1)2^{2f(\rho^2)}$. These estimates are summarized in Table 1.

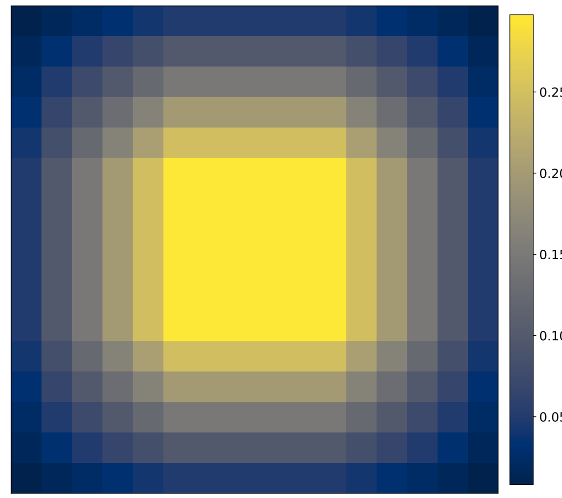

Figure 8: Density grid of the 6×6 block with $\rho = 5$. The color bar indicates values of Ψ. All points with $\Psi = \frac{36}{121} \approx 0.3$ are 1-components (the center yellow square).

	Interaction	**Dependency**
Vertices	$2^{f(\rho^2)}$	$(3\rho + 1)^2 2^{f(\rho^2)}$
Edges	$\gamma 2^{2f(\rho^2)}$	$\frac{\gamma}{2}(5\rho^2 + 6\rho + 1)2^{2f(\rho^2)}$

Table 1: Table of scaling estimates for the interaction and dependency graphs of blocks in LtL.

We now have a template for how to proceed to other model universes. First, we identify a criterion for distinguishing unities, and thus also the physical boundaries of such unities. Second, we want to define processes as spatial-chemical operations that induce dependency relations between processes that are local in space and sequential in time. Finally, the main object we will want to define is the closed network of dependency relations between production processes, such that the components thus produced constitute the system as a distinct unity in space. From this network can be derived a cognitive description of the system.

RealLife

RealLife is a family of Euclidean automata of arbitrary dimension (Pivato, 2007). The set of all possible 2-dimensional universes in RealLife is defined as a set of functions $\mathbf{u} : \mathbb{R}^2 \to \mathcal{M}$. Neighbourhoods can be defined with any p-norm, where $p = \infty$ gives a square and $p = 2$ gives a circle: $\mathcal{N}_\rho^x = \{x - y : ||y||_p \leq \rho, \, y \in \mathbb{R}^2\}$. Here, $\rho \in \mathbb{R}^+$. Thus, the convolution of Equation (2) can now be defined in continuous space:

$$\kappa * \mathbf{u}(x) = \lambda \left[\mathcal{N}_\rho\right]^{-1} \int_{\mathcal{N}_\rho} \mathbf{u}(x - y) \, d\lambda[y] \qquad (4)$$

where λ is the 2-dimensional Lebesgue measure on $\mathbb{R}^2$. We also change the rule parameters to satisfy $b_0 < b_1$ (instead of $b_0 \leq b_1$, as in LtL). The update rule in RealLife is the same as the LtL rule, except we use Equation (4) instead of Equation (2).

Pivato has proven the existence of still-lifes in RealLife (2007, Proposition 3.3). Specifically, let $\mathcal{B}^0 \subset \mathbb{R}^2$ be a $\rho \times \rho$ square of 1-components. Then, if $\left(\lambda[\mathcal{B}^0]/\lambda[\mathcal{N}_\rho]\right) \in [s_0, b_0)$, $\mathbb{1}_{\mathcal{B}^0}$ is a still-life that generalizes the block. As before, we will identify $\mathcal{B}^0$ as the internal 1-components and the immediately surrounding ring of width ρ as the block's boundary (Figure 9).

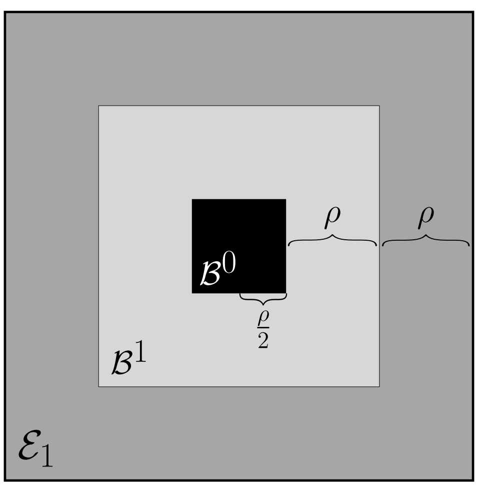

Figure 9: The block defined in RealLife with the max-norm. $\mathcal{B}^0$ is the region of internal 1-components, $\mathcal{B}^1$ is the boundary of 0-components, and $\mathcal{E}_1$ is the region of unspecified environmental components.

One of the scaling properties in LtL noted above is immediately apparent here: individual components are *irrelevant* to the dynamics. This is because the convolution is defined in terms of measure, and all countable sets have a Lebesgue measure of 0. As such, we will want to abstract over such variations as they are not really equivalent to the more interesting cases of structural degeneracy where variation does, in fact, have dynamical implications.

In defining a process, we can carry over much of what we did in LtL (Figure 10). The main complication introduced by the continuous case is the abstraction over countable variation. However, the central component of a given process cannot be included in those countable sets since the product would otherwise be ill-defined. We can notate subsets of $\mathcal{N}_\rho$ as belonging to the block or the environment, respectively, by a partition $\{\phi', \phi''\}$.

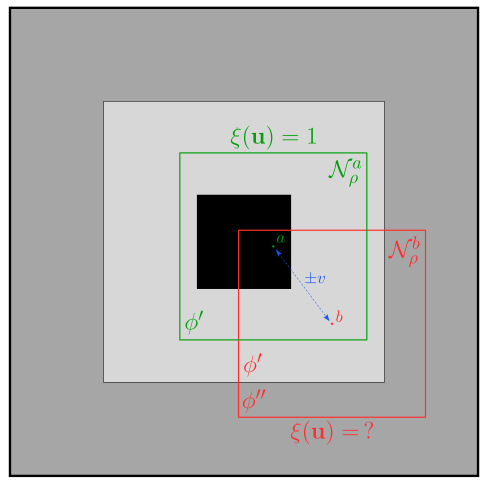

Figure 10: Processes and dependency in RealLife. The depicted dependency relation is between the canonical block configuration and itself. ϕ' denotes the region belonging to the block; ϕ'' denotes the region belonging to the environment; v is the dependency vector between the processes centered at a and b. The product $\xi(\mathbf{u})$ at b is potentially underdetermined, since $\lambda[\phi''] > 0$.

The dependency between processes, in contrast, is more difficult to deal with. The block, at least, makes matters simpler since we need not worry about translation between structures. A dependency link then becomes a spatial relation between processes. A vector between central components is sufficient for this (Figure 10). The direction of this vector depends on whether we are talking about one process *enabling* another, or one process being *dependent* on another. Both formulations turn out to give us the same information.

However, this is insufficient to define the *dependency manifold* as abstracted from the particular spatial embedding of the block. What we need is a way to define a space of processes in an essentially spatial manner without explicit dependence on the coordinates of the universe in which the block is embedded. One possibility is suggested by the similarity of adjacent processes: we could define a set of functions that transform processes in a way equivalent to a spatial shift of the neighbourhood (Figure 11). The main difficulty would lay in defining a constraint on these functions such that only valid processes would be included in the codomain. Moreover, these functions and their constraints may produce more than one viable configuration, further demanding a way to partition the codomain so as to separate these.

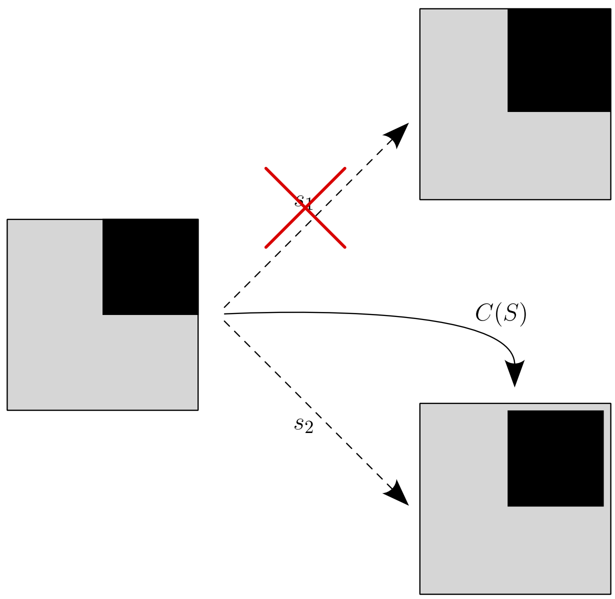

Figure 11: Schematic of functions and a constraint that together generate an abstract process space. S is a set of transformation functions s_1 and s_2. The process on the left is valid. Applying s_1 generates an invalid process. Applying s_2 generates a valid process. Thus, the constraint C on S selects s_2.

Once this is done, though, we can define dependency relations using vectors in this space, as opposed to the coordinate space of $\mathbb{R}^2$. Thus, the continuous generalization of the dependency graph of a block appears to be given by the neighbourhood relations in an abstract process space. However, it is not obvious how we should define the dependencies between different structures: do they exist in the same process space? If so, then how do we differentiate the dependencies a given process has on different structures? If not, then it appears we need an additional function to map processes to other structures, dependent on the existence of viable transitions between them.

Continuity further introduces another complication into structural degeneracy: the set of viable block structures is potentially uncountable. Interestingly, though, the transitions between these would still be discontinuous since time remains discrete in RealLife. This suggests that the *interaction manifold*, generalizing the interaction graph, exists in a continuous space of structures with discrete mappings between them. Again, it is not obvious how we should define this space in general, though it may be possible to parameterize certain regions in a more tractable manner. For instance, we could define a class of block configurations with a ring of some width embedded in its boundary. Then we could analyse what transitions are possible — if any — within this class.

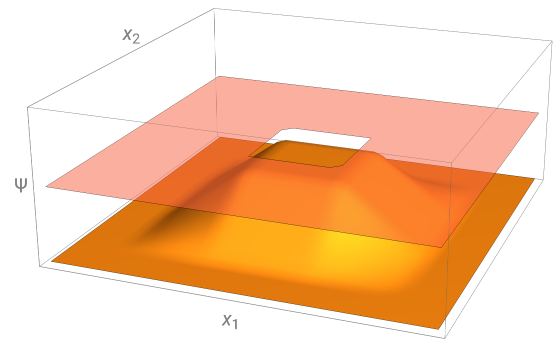

Figure 12: Smooth density manifold generated by Ψ for a block in an empty environment in RealLife. In orange is the manifold, and in pink is a constraint over the boundary. x_1 and x_2 are the coordinate axes of $\mathcal{B}^{(0,1)}$. A perturbation can increase the density over the boundary. If any point on the manifold reaches the constraint, the transition induced on the structure changes.

To carry out such an analysis, we again shift our focus to local density, now forming a manifold (Figure 12). As with the block in LtL, we can define a constraint that induces specific transitions (the pink surface in Figure 12). Thus, it is hoped, we can begin to get an analytical grasp on the smaller sub-problems such that we can develop the tools we will need to address more general problems.

Discussion

The goal of this paper has been to sketch the landscape of problems in using Larger than Life and RealLife as model universes for studying the autopoietic organization of emergent individuals. I showed how the concepts and definitions used by Beer in the Game of Life cellular automaton could be extended in a fairly straightforward way to LtL. I also pointed out certain features of how LtL scales and proposed density grids and constraints as useful analytic constructs to cope with the growth of objects scaled with the neighbourhood. Finally, I provided an outline of the problems and possibilities introduced in RealLife.

The main challenges now apparent in using RealLife are in defining abstract process and structure spaces independent of the coordinate space in which an entity is embedded. And even when this is done, there still remain questions as to how we should relate structural degeneracies and translations in these spaces.

The density manifold appears to be a smooth object, and thus it would be useful if we could differentiate on it. For instance, we may want to find local maxima that are tangent to the constraint in order to determine the set of perturbations that induce a given transition. However, since Ψ is defined as a convolution of two indicator functions, its derivative is always zero, or else undefined. Thus, if we are to define — or else approximate — the density manifold in a differentiable manner, different mathematical tools will be needed.

Moreover, the existence of oscillators and bugs has yet to be proven in RealLife. It may also be the case that, at the continuum limit, these structures are no longer perfectly periodic (Pivato, 2007, pp. 67). Then it becomes imperative to analyse these structures with a modified closure condition, but to ensure that we do so in a theoretically principled manner.

The questions pertaining to autopoiesis that have concerned us here are only a small fraction of a far larger space of ideas. For instance, we could also investigate the relation between *operational* and *physical* boundaries, that have here been conflated (Virgo et al., 2011). Such an investigation would deal intimately with the non-arbitrarity of the conditions of observation; we would also need to understand how the epistemological concepts used in the biology of cognition (Varela, 1979; Maturana, 2002) translate to these model universes. Beer has previously investigated the origins and structural coupling of gliders in GoL (Beer, 2020a,c), so there is already some groundwork on which these ideas could be further pursued. One could also compare the toy chemistries of LtL and RealLife with various formalisms in theoretical chemistry, as has previously been done for Lenia and reaction-diffusion systems (Kojima and Ikegami, 2023).

In discussing structural degeneracy, I employed the concept of *viability*. This notion has a rich history with autopoiesis (Varela, 1979; Bourgine and Varela, 1992), and much independent theoretical development (Ashby, 1960; Aubin, 1991; McShaffrey and Beer, 2023). With respect to living systems, viability theory aims to delineate the existence and non-existence of the system as a distinct entity — the boundary between life and death. Thus, an understanding of viability provides a constraint that bounds the conditions under which a system can prolong its existence through its continued operation. Such a constraint can be derived directly from the organization of an emergent individual (Beer et al., 2024), but under certain conditions, we can get a sense of the viability constraint using density manifolds. Specifically, if we assume knowledge of an individual's viable structures, and that the individual does not move (does not gain environmental components), we can assign density constraints on a given viable structure such that each corresponds to a transition in the cognitive domain. Density constraints then provide a local picture of viability by together specifying what perturbations cannot induce any viable transition from a given structure. For example, if we consider the canonical block to be the only viable structure, then Figure 12 depicts the boundary between conditions under which the block does or does not persist, depending on whether the manifold is below or above the constraint, respectively. Hence, density manifolds serve as a practical alternative to deriving constraints directly, even though they are theoretically subordinate to the organization (the viability of structures is determined by the organization) and, in their current form, do not account for moving individuals.

Finally, Lenia offers an interesting contrast to RealLife in being continuous in both space *and* time (Chan, 2019). This is significant because the mathematical tools needed to satisfactorily analyse such systems from an autopoietic perspective, at the moment, would have to be invented from scratch. Thus, using RealLife as a stepping stone to the fully continuous case creates a bridge between structures in the entirely discrete universe of GoL and entirely continuous universes such as Lenia.

Acknowledgements

I would like to thank Connor McShaffrey and Randall Beer for their helpful discussions and comments.

References

Adamatzky, A. (2010). *Game of Life Cellular Automata*. Springer-Verlag.

Ashby, W. R. (1960). *Design For a Brain: The Origin of Adaptive Behaviour*. Springer Science & Business Media.

Aubin, J. P. (1991). *Viability Theory*. Birhäuser.

Beer, R. D. (2004). Autopoiesis and cognition in the Game of Life. *Artificial Life*, 10(3):309–326.

Beer, R. D. (2014). The cognitive domain of a glider in the Game of Life. *Artificial life*, 20(2):183–206.

Beer, R. D. (2015). Characterizing autopoiesis in the Game of Life. *Artificial life*, 21(1):1–19.

Beer, R. D. (2017). Computing aggregate properties of preimages for 2D cellular automata. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 27(11):111104.

Beer, R. D. (2020a). Bittorio revisited: structural coupling in the Game of Life. *Adaptive Behavior*, 28(4):197–212.

Beer, R. D. (2020b). An integrated perspective on the constitutive and interactive dimensions of autonomy. In Bongard, J., Lovato, J., Herbert-Dufrésne, L., Dasari, R., and Soros, L., editors, *ALIFE 2020: The 2020 Conference on Artificial Life*, pages 202–209. MIT Press.

Beer, R. D. (2020c). An investigation into the origin of autopoiesis. *Artificial Life*, 26(1):5–22.

Beer, R. D., McShaffrey, C., and Gaul, T. M. (2024). Deriving the intrinsic viability constraint of an emergent individual from first principles. In *Proceedings of the 2024 Artificial Life Conference*. To appear.

Bourgine, P. and Varela, F. J. (1992). Towards a practice of autonomous systems. In Bourgine, P. and Varela, F. J., editors, *Toward a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life*, pages xi–xvii. MIT Press.

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

Di Paolo, E. A. (2005). Autopoiesis, adaptivity, teleology, agency. *Phenomenology and the Cognitive Sciences*, 4(4):429–452.

Evans, K. M. (2001). Larger than Life: Digital creatures in a family of two-dimensional cellular automata. *Discrete Mathematics and Theoretical Computer Science*, 177:192.

Evans, K. M. (2003). Larger than Life: Threshold-range scaling of life's coherent structures. *Physica D: Nonlinear Phenomena*, 183(1-2):45–67.

Kojima, H. and Ikegami, T. (2023). Implementation of Lenia as a reaction-diffusion system. In Iizuka, H., Suzuki, K., Uno, R., Damiano, L., Spychala, N., Aguilera, M., Izquierdo, E., Suzuki, R., and Baltieri, M., editors, *ALIFE 2023: Ghost in the Machine: Proceedings of the 2023 Artificial Life Conference*, pages 43–49. MIT Press.

Maturana, H. (2002). Autopoiesis, structural coupling and cognition: a history of these and other notions in the biology of cognition. *Cybernetics & Human Knowing*, 9(3-4):5–34.

Maturana, H. R. and Varela, F. J. (1980). *Autopoiesis and Cognition: The Realization of the Living*. D. Riedel Publishing.

Maturana, H. R. and Varela, F. J. (1987). *The Tree of Knowledge: The Biological Roots of Human Understanding*. Boston: Shambala.

McMullin, B. (2004). Thirty years of computational autopoiesis: A review. *Artificial life*, 10(3):277–295.

McShaffrey, C. and Beer, R. D. (2023). Decomposing viability space. In Iizuka, H., Suzuki, K., Uno, R., Damiano, L., Spychala, N., Aguilera, M., Izquierdo, E., Suzuki, R., and Baltieri, M., editors, *ALIFE 2023: Ghost in the Machine: Proceedings of the 2023 Artificial Life Conference*, pages 357–364. MIT Press.

Pivato, M. (2007). RealLife: The continuum limit of Larger than Life cellular automata. *Theoretical Computer Science*, 372(1):46–68.

Varela, F. G., Maturana, H. R., and Uribe, R. (1974). Autopoiesis: The organization of living systems, its characterization and a model. *Biosystems*, 5(4):187–196.

Varela, F. J. (1979). *Principles of Biological Autonomy*. New York: North Holland.

Virgo, N., Egbert, M. D., and Froese, T. (2011). The role of the spatial boundary in autopoiesis. In Kampis, G., Karsai, I., and Szathmáry, E., editors, *Advances in Artificial Life. Darwin Meets von Neumann*, pages 240–247. Springer Berlin Heidelberg.

von Uexküll, J. (1992). A stroll through the worlds of animals and men: A picture book of invisible worlds. *Semiotica*, 89(4):319–391.