

Deriving The Intrinsic Viability Constraint of an Emergent Individual from First Principles

Randall D. Beer, Connor McShaffrey, and Thomas M. Gaul

Cognitive Science Program
Indiana University
rdbeer@iu.edu

Abstract

The notion of a viability constraint that determines the range of conditions under which a biological individual can survive plays a central role in work in artificial life and theoretical biology. However, while there has been considerable attention paid to the case where this constraint is defined externally, very little work has been done on the more natural case where this constraint arises intrinsically from the operational closure of the individual itself. Using a glider in the Game of Life as a toy model, we show how to systematically derive the intrinsic viability constraint of an emergent individual from its closed network of constitutive process interdependencies.

Introduction

One of the most fundamental features of any living system is the boundary between life and death. On one side of this boundary, the system exists as a distinct entity, capable of engaging its environment as an autonomous individual. On the other side of this boundary, there is no longer any system to discuss; inquiries into its state or goals or behavior become meaningless. This boundary is often envisioned as a kind of constraint – a *viability* constraint – on the system's state variables which delineates its domain of existence.

It is useful to distinguish two broad classes of viability constraint: extrinsic and intrinsic. An *extrinsic* viability constraint is one imposed from the outside by some external agency. For example, a protocell model may be considered to have “died” if the concentration of some essential chemical falls below some lower limit or its osmotic pressure exceeds some upper limit. While this might seem entirely artificial and arbitrary, for pragmatic reasons it is by far the most common approach in theoretical and computational biology and artificial life. The cybernetician W. Ross Ashby, who was one of the first to make explicit the concept of a viability constraint, conceived of it as a compact subset of a dynamical system's state space involving the system's essential variables (Ashby, 1960). Similar notions can be found in control theory (Aubin, 1991), adaptive behavior (Beer, 1995), health (Voit, 2009), artificial life (Barandarian & Egbert, 2013), and computational cell biology (Ghaffarizadeh et al., 2018). Recent work in this area has focused on formally characterizing the global structure of

viability outcomes when a dynamical system possesses such constraints (McShaffrey & Beer, 2023).

In contrast, an *intrinsic* viability constraint arises from within a living system itself rather than being externally imposed. What does this even mean? A key requirement for a system to possess an intrinsic viability constraint is that it must be systemically precarious, which can only happen if its identity is emergent from some underlying substrate to which it can subsequently decay when its integrity as an individual is lost (Beer & Di Paolo, 2023). In this case, the viability limits arise from the organization and operation of the constituent components of the system itself. Whereas examples of extrinsic viability constraints are ubiquitous, despite the fact that intrinsic viability is arguably the more natural and fundamental concept, there are very few examples of such emergent individuals in the literature (Varela, Maturana & Uribe, 1974; Ono & Ikegami, 2000; McMullin, 2004; Agmon et al., 2016). Even for these few examples, no explicit construction of the resulting intrinsic viability constraint has ever been carried out.

Developing a theory of viability has primarily been an independent endeavor, but the subject also has a long-standing history with Maturana and Varela's formulation of the organization of living systems, *autopoiesis* (Maturana & Varela, 1980). Varela (1979) suggested an equivalence between the two concepts. Later, Bourguin and Varela (1992) argued for a close relation instead, with viability specifying the bounds within which the organization's operation would continue. Previously, these connections were presented conceptually, with the assumption that the viability domain was already known to the observer.

In this paper, we show how to derive the intrinsic viability constraint of an emergent individual from first principles using a glider in the Game of Life cellular automaton as an example. We first present the organization of a glider as a closed process dependency graph. We then reconstruct the spatial embedding of this graph from its dependency structure. Next, we show how the spatially embedded organization determines a glider's graph of identity-preserving interactions. Finally, we extract and visualize its intrinsic viability constraint from this interaction graph.

Process Dependency Network

Our derivation of an intrinsic viability constraint builds on our previous investigations of a glider in Conway's Game of Life (GoL) cellular automaton as a model of an emergent individual (Beer, 2004). In this model, the underlying GoL update rules determine a kind of physics (called the *Conway physics*), which supports a kind of spatial chemistry, which in turn supports the emergence of closed networks of processes whose self-generating and self-bounding nature can be analyzed from the perspective of Maturana and Varela's (1980) notion of autopoiesis.

In this framework, a process is a rule for transforming a local spatial configuration of reactants into a product (Figure 1A). Processes fall into four main classes: production (blue), destruction (red), 1-maintenance (black), and 0-maintenance (white) according to their action on the cell at their center, an action determined by the states of the immediately surrounding cells according to the Conway physics. By tracking how the product of each process in turn enables other processes for an isolated glider in an otherwise empty universe, one can extract the closed process dependency network that constitutes what we later called the *vacuum glider organization* (Beer, 2015).

In order to generalize this description of a glider's organization to nonempty universes, we subsequently introduced the notion of a *partial process* (Beer, 2020a). Briefly, a partial process is a regular process some of whose triggering cells have been replaced by an "unknown" state in order to represent entire classes of actual processes (Figure 1B). Partial processes are deployed in situations, such as those that occur along the interface of a glider with its environment, in which the actual states of some cells are unknown until the glider is embedded in a particular environment. In this case, the partial boundary processes take on different roles depending on the context (indicated by the triangular split of the central cells in Figure 1B). For example, the partial process illustrated to the left of Figure 1B represents all processes whose central cell is currently OFF and whose four cells outlined in black take on the given values. Depending upon the sum of the remaining unspecified cells (yellow), this process could be either a 0-maintenance or a production process. The right of Figure 1B shows the analogous situation when the central cell is currently ON rather than OFF.

By replacing the boundary processes in the vacuum glider organization with the corresponding partial processes and tracing out the consequences of different environmental contexts on these processes, one obtains the *generalized glider organization* (Beer, 2020a). This organization, which has not previously been presented in its entirety, is shown in Figure 2. This graph includes a total of 352 processes and 5 912 dependency links. Note that the same process may appear multiple times but still have distinct identities because their dependencies differ, a consequence of the spatiality of the chemistry supported by the Conway physics. Note also that 32 of these processes (circled in red) receive only incoming dependency links and thus do not participate in the operational closure of the glider; they represent processes associated with "waste" components that the glider sheds during its operation.

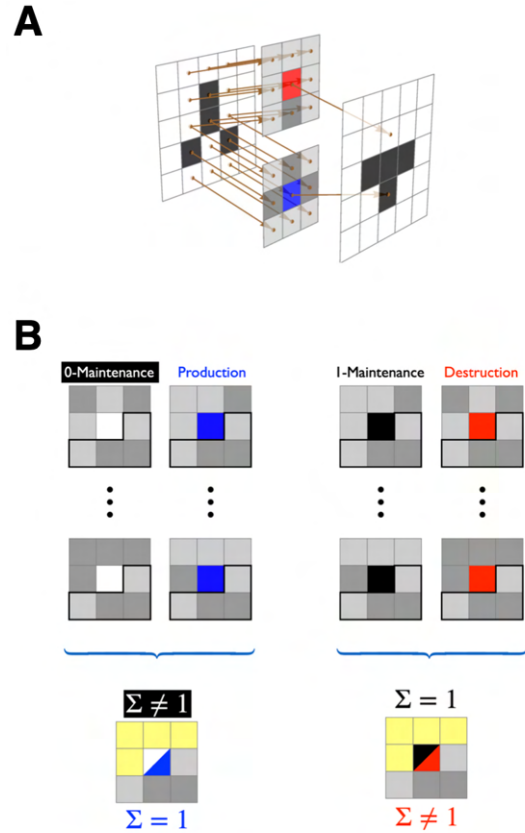


Figure 1: Processes in the Game of Life. (A) The action of a production process (blue central square) and a destruction process (red central square) on a configuration of components, with brown links indicating dependencies between components and processes. (B) A partial process can represent a set of complete processes. All processes shown contain the same cell states in the L-shaped region outlined in black, with the other cell states varying. The left partial process represents all such processes triggered by a 3 × 3 configuration whose central cell is currently OFF. Depending on the sum Σ of the values of the unspecified yellow cells, the partial process can become either a 0-maintenance or production process, hence the split white/blue central cell. The right partial process represents the analogous situation when the central cell is currently ON, triggering either a 1-maintenance or destruction process depending on the sum of the unspecified yellow cells.

Local Spatial Embedding

A glider's organization is a topological entity; it depends only on the interdependencies between a glider's constituent processes. However, as we will show in this section, the spatial structure of a glider can be fully recovered from this dependency network, at least up to embedding symmetries.




Figure 2: The (generalized) glider organization expressed as a graph of dependency links between the products and enabling conditions of a closed set of processes. The processes circled in red receive only incoming dependency links so they are not part of the closure; they represent processes involving “waste” components that the glider sheds during its operation.

The basic idea is as follows. Each process describes an “atom” of geometry through the relative spatial relations of its nine constituent cells that together constrain what configurations of adjacent processes are possible. Furthermore, the dependency links between the cells of different processes similarly determine a set of geometrical constraints that any embedding of the processes must satisfy. For example, given a precursor process, the set of processes it enables must be adjacent and thus arranged in such a way that each process has its neighborhood constraints satisfied. The problem of

assembling a spatial embedding from these constraints thus becomes a kind of jigsaw puzzle that can be solved algorithmically.

The algorithm has two main steps. First, we create a set of local spatial embeddings for each process by splitting the processes enabled by its outgoing dependency links into subsets that agree on both the product and the local spatial arrangement (Figure 3A). Second, we repeatedly merge pairs of local embeddings that contain the same processes in such a way as to respect the geometry of each until no further merges




Figure 3: An illustration of each of the two key steps involved in spatially embedding a process dependency network. (A) Each subset of a process’s dependency links that agree on a product and a consistent spatial arrangement imply a partial embedding of its subsequent processes. (B) Two partial embeddings that contain overlapping processes (highlighted in red) can be merged.

are possible (Figure 3B). Note that if the “waste” processes are removed from the process dependency graph (Figure 2), they need to be reintroduced here in an additional step by completing any dangling dependencies in order for this algorithm to fully recover the spatial embedding.

The result of applying this algorithm to the generalized glider organization presented in the previous section is shown in Figure 4. This result completes the partial visualization of the local spatial embedding of the glider organization that was shown in (Beer, 2020a) and used as the starting point for the work described therein. Note that the spatial embedding is only semi-local; the relative spatial relations within each of the sixteen distinct process sets are made explicit, but not the spatial relations between those distinct process sets.

Interaction Graph

The spatially embedded process dependency graph fully describes a glider’s constitution, characterizing the way that its components and their interactions give rise to and maintain the coherent identity of this emergent individual. However, one can also describe the glider as a unitary whole interacting with its environment. Maturana termed the latter description a “simple

unity” and the former description a “composite unity” (Maturana & Varela, 1980). For a simple unity, one is most interested in characterizing the possible interactions with its environment in which it can engage. This “cognitive domain” (Maturana & Varela, 1980) describes all possible structural configurations of a glider and the sets of environmental perturbations that trigger transitions between them.

One way to compute a glider’s interaction graph is to exhaustively enumerate a glider’s response to each possible environmental perturbation (Beer, 2014). However, in more recent work, we showed that it is also possible to derive a glider’s interaction graph directly from its spatially embedded process dependency network (Beer, 2020a). The derivation proceeds in three steps (Figure 5).

The first step is a simple relabeling of the process dependency links (Figure 5A). All dependency links between any two pairs of glider process embeddings are colored the same way, with the chosen color scheme deployed consistently so as to respect the symmetries of the spatially embedded dependency network.

The second step involves an abstraction: each of the sixteen glider process embeddings is replaced by the corresponding




Figure 4: The local spatial embedding of the process dependency graph shown in Figure 2 as produced by the algorithm described in the main text and outlined in Figure 3.

configuration of components that trigger them (Figure 5B). After this step, we see that each process embedding corresponds to a different glider configuration, corresponding to the four different orientations and two different chiralities of each of the two possible glider forms.

The third and final step is to collapse the bundles of process dependency links between all pairs of glider configurations to a single edge (Figure 5C). The resulting graph is identical to the glider interaction graph that has previously been reported (Beer, 2014). Here each node represents a possible glider configuration and each color of link represents a class of identity-preserving environmental perturbations that trigger the corresponding transition. We also showed how to explicitly

calculate every environment configuration belonging to each of these six classes of nondestructive perturbations (Beer, 2014). Of course, this interaction graph only describes a glider's response to a given environmental perturbation. In order to completely determine the possible lives that a glider can live, one must also characterize the environment's responses to glider perturbations, as well as how these two interaction graphs interlock via structural coupling (Beer, 2020b).

Note that, strictly speaking, Figure 5C shows only the translation symmetry reduced glider interaction graph. In order to generate the full interaction graph, one must unroll this graph across all possible locations in a given GoL universe (Beer, 2020b).




Figure 5: Steps in the derivation of a glider's interaction graph from the local spatial embedding of the glider process dependency network shown in Figure 4. (A) Dependency links between any pair of process sets are colored in a way consistent with the symmetries of Figure 4. (B) Each process set is replaced by the spatial arrangement of components that trigger them. (C) The bundle of like-colored dependency links between any two glider configurations can be reduced to a single transition link, reducing a glider's spatially embedded process dependency network to its interaction graph.

Viability Constraint

The glider interaction graph shown in Figure 5C gives us every environmental perturbation that a glider can survive. Its complement gives us every environmental perturbation that destroys a glider. Together, this information can be used to derive the glider's intrinsic viability constraint.

The resulting viability constraint is illustrated schematically in Figure 6. The space $\mathcal{G}^{(1,2)}$ in which this constraint is visualized is the space of all sets of GoL cells that could contain a glider at its center, surrounded by that glider's

immediate environment (called the 1-environment), surrounded by its secondary environment (called the 2-environment). It is necessary to include the 2-environment because when a glider moves, it occupies some of the cells that had previously constituted its 1-environment, whose states are co-determined by its 2-environment. Since there are 78 GoL cells in each element, $\mathcal{G}^{(1,2)}$ contains 2^{78} elements.

The glider interaction graph divides $\mathcal{G}^{(1,2)}$ into three subsets. The union $\text{int}(\mathcal{V}) \cup \partial\mathcal{V}$ contains all configurations possessing a central glider, with the viability interior $\text{int}(\mathcal{V})$ (light gray) including all configurations whose central glider will persist to the next step and the viability boundary $\partial\mathcal{V}$ (dark




Figure 6: A schematic illustration of the structure of a glider's derived intrinsic viability constraint. The viability constraint \mathcal{V} (gray) contains all configurations that possess a central glider, whereas its complement $\bar{\mathcal{V}}$ contains all configurations that do not. \mathcal{V} can be further divided into an interior $\text{int}(\mathcal{V})$ (light gray) and a boundary $\partial\mathcal{V}$ (dark gray), with $\mathcal{V} = \text{int}(\mathcal{V}) \cup \partial\mathcal{V}$. The region $\text{int}(\mathcal{V})$ contains all configurations that will remain in \mathcal{V} after the next transition, whereas $\partial\mathcal{V}$ contains all configurations that will leave \mathcal{V} (indicated by the red arrows) after the next transition. A few example configurations from each region are shown.

gray) containing all configurations whose central glider will perish at the next step (in this discrete space, $\text{int}(\mathcal{V})$ and $\partial\mathcal{V}$ are distinct sets). We will refer to this union as the glider's viability region \mathcal{V} (gray). The complement of \mathcal{V} , $\bar{\mathcal{V}}$, includes all elements of $\mathcal{G}^{(1,2)}$ that do not possess a central glider. Calculations show that $\text{int}(\mathcal{V})$ and $\partial\mathcal{V}$ contain 183 360 743 800 832 and 143 876 619 822 432 256 elements, respectively, with $\bar{\mathcal{V}}$ containing 302 231 310 843 676 727 443 456 elements.

Note that the Conway physics induces a dynamics on $\mathcal{G}^{(1,2)}$ that is only partly represented in Figure 6. As described by the glider interaction graph, elements in $\text{int}(\mathcal{V})$ transform into other elements in $\text{int}(\mathcal{V})$ through nondestructive perturbations. In addition, because these configurations are embedded in larger environments that can modify a glider's 2-environment, elements in $\text{int}(\mathcal{V})$ can transform into elements of $\partial\mathcal{V}$. In turn, elements of $\partial\mathcal{V}$ inexorably transform into elements of $\bar{\mathcal{V}}$ as the central glider disintegrates, as indicated by the red arrows. Finally, elements of $\bar{\mathcal{V}}$ can transform into elements of $\text{int}(\mathcal{V})$ or $\partial\mathcal{V}$ as new gliders are spontaneously created from glider precursor configurations. Interestingly, these $\mathcal{V} \rightarrow \mathcal{V}$, $\mathcal{V} \rightarrow \bar{\mathcal{V}}$, and $\bar{\mathcal{V}} \rightarrow \mathcal{V}$ transitions correspond precisely to the glider persistence, destruction, and creation operators described in (Beer, 2020c).

Studying the process dependencies and the resulting interaction graph reveals a more nuanced picture of essential variables than is typically presented in the extrinsic viability literature, where the set is pre-determined and fixed. Here, the components that participate in a glider's continued viability depend on the particular transition in its cognitive domain. For example, a glider's autopoiesis will not rely on boundary cells that it "sheds" in its movement, but the direction of movement depends on the joint state of the glider and its (1,2)-environment. By this framing, the particular set of "essential variables" at play in any given instant is transition-dependent and includes environmental elements.

Having fully characterized the glider's generalized organization and viability constraint, we can now draw more explicit connections between the nature of autopoiesis and viability. For example, autopoiesis is characterized by the closure of processes that results in the continuation of the glider. Viability, on the other hand, is the domain of configurations where the glider can be distinguished, including those that will immediately lead to its death. Thus, while viable transformations of the glider's structure are subservient to autopoiesis, the viability constraint specifies the physical limits at which the organization will cease to exist. This makes

viability a function of both the organization and the structural characteristics of the glider and its immediate environment.

Discussion

The viability of a biological organism is determined not from the outside, but from within, based on the continued integrity of the constitutive network of processes that underlie its operation. In an attempt to characterize how such an intrinsic viability constraint arises from first principles, we have examined a glider in the Game of Life as a toy model of an emergent individual. Specifically, starting from a glider's organization as a closed network of process dependencies, we have shown how to incrementally derive (1) its local spatial embedding, (2) its interaction graph, and, ultimately, (3) the intrinsic viability constraint implicit in that organization. We have also begun to examine the nontrivial internal structure of this intrinsic viability constraint. To our knowledge, this is the first time that an intrinsic viability constraint has been derived and characterized for any emergent individual.

Having a genuine instance of a derived viability constraint allows the rare opportunity to reflect on some broader themes surrounding the subject. For example, one characterization of adaptivity discusses organisms as distinguishing between dynamics that will result in approaching or moving away from the boundary of their viability region (Di Paolo, 2005; 2009). By this framing, agents sense their viability constraints and respond to them via structural changes. Contrary to this, we have shown that the viability region of the glider *originates* from the nature of its potential structural transitions. While the glider is sensitive to its own state and that of its immediate environment, it cannot sense the 2-environment cells that participate in the boundary of the viability region. This raises the distinction between the glider having access to its set of potential essential variables versus the viability boundary itself, consistent with previous work demonstrating that adaptivity does not require an explicit sensing of viability's limits (Agmon, Gates & Beer, 2016). It has also been shown that an agent can engage in viable behavior without sensing essential variables directly so long as it can sense an adequate proxy, although this behavior may be more fragile (Agmon & Beer, 2014; Egbert & Pérez-Mercader, 2016).

There are also comparisons to be made between the anatomy of the glider's intrinsically generated viability region and the viability constraints that we extrinsically define. In extrinsic viability, the constraints are defined entirely separately of the equations that govern the dynamics. The result of this is that the states populating the boundary can be characterized by change vectors that push the system into the terminal region, pull it back into the interior, or run tangent. In the intrinsic case, $\mathcal{G}^{(1,2)}$ configurations are assigned to the boundary *because* they result in terminal transitions, and therefore border it. This difference is especially notable given how important the different vector types are in the global analysis of extrinsic viability (McShaffrey & Beer, 2023).

While it may seem as though the intrinsic derivation escapes some of the artificial decisions of the extrinsic case, it

is worth noting that, even here, there is a sense in which the structure we derive is still partially observer dependent. Currently, we say that a GoL unity has disintegrated if its organization is disrupted to any degree. If we instead decided that the unity survived if it managed to reconstitute immediately following any disruption, this would change the interior structure and boundary of the viability region.

Despite the differences in their properties, co-developing theories of intrinsic and extrinsic viability allows us to draw on the strengths of both. Since models of intrinsic viability are the closest to the naturalistic case, they can teach us about the foundation of viability constraints, and in some cases it may even be possible to extract them to utilize in extrinsic models. Extrinsic viability, on the other hand, gives us a framework to approach problems where modeling a full, emergent individual is not possible. There is also the possibility of using classification schemes in extrinsic viability to further direct research on intrinsic viability. For example, viability space decomposition carves up extrinsic models into asymptotically and transiently viable regions based on whether the agent will persist or die on a finite time horizon (McShaffrey & Beer, 2023). Merging the glider with an environment gives an autonomous dynamical system where similar global classifications should be possible, building on previous work on structural coupling in GoL (Beer, 2020b).

Looking forward, it would be easy to apply the procedure outlined in this paper to derive the intrinsic viability constraints for other emergent individuals in the Game of Life. Extension to other discrete cellular automata and discrete spatial artificial chemistries should also be relatively straightforward. However, generalizing our approach to the continuous case, such as reaction-diffusion systems (Froese, Virgo & Ikegami, 2011), continuous spatial artificial chemistries (Dittrich, Ziegler & Banzhaf, 2001), or continuous cellular automata such as Lenia (Chan, 2019), raises a number of subtle issues, mostly deriving from spatiality. For example, a continuous emergent individual is constituted by a *manifold* of processes rather than a discrete grid, and its process dependencies also form a manifold rather than a discrete graph. It is not at all clear what the continuous generalization of a discrete process dependency network should be and therefore it is not yet clear how to carry out the steps of the procedure outlined in this paper for the case of a continuous emergent individual. An interesting approach might be to investigate the Larger Than Life cellular automaton (Evans, 2003) and its continuum limit RealLife (Pivato, 2007) as an incremental path from the discrete to the continuous case (Gaul, 2024).

Acknowledgements

This material is based upon work supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. 2240777. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

References

- Agmon, E. and Beer, R.D. (2014). The evolution and analysis of action switching in embodied agents. *Adaptive Behavior* **22**:3-20.
- Agmon, E., Gates, A.J., Churavy, V. and Beer, R.D. (2016). Exploring the space of viable configurations in a model of metabolism-boundary co-construction. *Artificial Life* **22**:153-171.
- Agmon, E., Gates, A.J., and Beer, R.D. (2016). The structure of ontogenies in a model protocell. *Artificial Life* **22**:499-517.
- Ashby, W.R. (1960). *Design for a Brain: The Origin of Adaptive Behavior* (Second Edition). Springer Science & Business Media.
- Aubin, J.P. (1991). *Viability Theory*. Birkhäuser.
- Barandarian, X.E. and Egbert, M.D. (2013). Norm-establishing and norm-following in autonomous agency. *Artificial Life* **20**:5-28.
- Beer, R. D. (1995). A dynamical systems perspective on agent-environment interaction. *Artificial Intelligence* **72**:173-215.
- Beer, R.D. (2004). Autopoiesis and cognition in the Game of Life. *Artificial Life* **10**:309-326.
- Beer, R.D. (2014). The cognitive domain of a glider in the Game of Life. *Artificial Life* **20**:183-206.
- Beer, R.D. (2015). Characterizing autopoiesis in the Game of Life. *Artificial Life* **21**:1-19.
- Beer, R.D. (2020a). An integrated perspective on the constitutive and interactive dimensions of autonomy. In J. Bongard et al. (Eds), *Proceedings of the 2020 Conference on Artificial Life* (pp. 202-209). MIT Press.
- Beer, R.D. (2020b). Bittorio revisited: Structural coupling in the Game of Life. *Adaptive Behavior* **28**:197-212.
- Beer, R.D. (2020c). An investigation into the origin of autopoiesis. *Artificial Life* **26**:5-22.
- Beer, R.D. and Di Paolo, E.A. (2023). The theoretical foundations of enaction: Precariousness. *BioSystems* **223**: 104823.
- Bourgine, P. and Varela, F.J. (1992). Towards a practice of autonomous systems. In Bourguine, P. and Varela, F.J. (Eds.), *Towards a Practice of Autonomous Systems* (pp. xi-xvii). MIT Press.
- Chan, B.W.-C.. (2019). Lenia: Biology of artificial life. *Complex Systems* **28**:251-286.
- Di Paolo, E.A. (2005). Autopoiesis, adaptivity, teleology, agency. *Phenomenology and the Cognitive Sciences* **4**:429-452.
- Di Paolo, E.A. (2009). Extended life. *Topoi* **28**:9-21.
- Dittrich, P., Ziegler, J., and Banzhaf, W. (2001). Artificial chemistries – A review. *Artificial Life* **7**:225-275.
- Egbert, M.D. and Pérez -Mercador, J. (2016). Adapting to adaptations: Behavioural strategies that are robust to mutations and other organizational-transformations. *Scientific Reports* **6**:18963.
- Evans, K.M. (2003). Larger than Life: Threshold range scaling of Life's coherent structures. *Physica D* **183**:45-67.
- Froese, T., Virgo, N. and Ikegami, T. (2011). Life as a process of open-ended becoming: Analysis of a minimal model. In T. Lenaerts et al. (Eds.), *Advances in Artificial Life 11 (ECAL 2011)*, (pp. 250-257). MIT Press.
- Gaul, T.M. (2024). Autopoiesis in RealLife euclidean automata. To appear in *Proceedings of ALife 2024*.
- Ghaffarizadeh, A., Heiland, R., Friedman, S.H., Mumenthaler, S.M. and Macklin, P. (2018). PhysiCell: An open source physics-based cell simulator for 3-D multicellular systems. *PLoS Computational Biology* **14**:e1005991.
- Maturana, H.R. and Varela, F.J. (1980). *Autopoiesis and Cognition: The Realization of the Living*. D. Reidel Publishing.
- McMullin, B. (2004). Thirty years of computational autopoiesis: A review. *Artificial Life* **10**:277-295.
- McShaffrey, C. and Beer, R.D. (2023). Decomposing viability space. In Iizuka, H., et al. (Eds.), *Proceedings of the Artificial Life Conference 2023* (pp. 357-364). MIT Press.
- Ono, N. and Ikegami, T. (2000). Self-maintenance and self-reproduction in an abstract cell model. *J. Theoretical Biology* **206**:243-253.
- Pivato, M. (2007). RealLife: The continuum limit of Larger Than Life cellular automata. *Theoretical Computer Science* **372**:46-68.
- Varela, F.J. (1979). *Principles of Biological Autonomy*. North Holland.
- Varela, F.G., Maturana, H.R., and Uribe, R. (1974). Autopiesis: The organization of living systems, its characterization and a model. *BioSystems* **5**:187-196.
- Voit, E.O. (2009). A systems-theoretic framework for health and disease: Inflammation and preconditioning from an abstract modeling point of view. *Mathematical Biosciences* **217**:11-18.