A Saucerful of Secrets: Open-Ended Organizational Closure in the Game of Life

Fernando Rodriguez and Phil Husbands

AI Research Group, Department of Informatics, University of Sussex
f.rodriguez-vergara@sussex.ac.uk

Abstract

Organizational closure and Open-endedness are both central, widely explored concepts within Artificial Life, although still requiring further theoretical development. In this work, we present a novel approach for characterizing organizational closure in terms of multidimensional mappings and, by building upon it, propose a principled conceptual connection to open-endedness, leading to the notion of open-ended organizational closure. Our fundamental claim is that both these properties are complementary and necessary for minimally adaptive autonomous systems, even prior to evolutionary considerations. As a proof of concept for this idea, we present an experimental analysis of a toy model setup, by examining the dynamics of the Game of Life cellular automaton. Finally we discuss theoretical implications, the limitations of our toy model scenario and suggest lines for further work.

Introduction

Years of research in artificial life, enactive cognitive science and related fields have given us good insights into some of the properties of cognition and its theoretical relations to life. Stemming from foundational work from Maturana and Varela (1973); Varela (1979); Varela et al. (1991), some understanding has been gained into how some particular types of self-sustaining systems are able to escape disintegration by drawing relevant distinctions from the world they inhabit, and how these manifest as autonomous behavior through ongoing system-environment co-determined changes (Froese and Di Paolo, 2011; Di Paolo et al., 2017). Such changes are underpinned, generally, by the recursive nature of a certain class of dynamics, denoted as autonomous; and particularly, by the specific intrinsic logic their autonomy determines, which is often referred to as their organization. That is, the abstract, overall collection of their possible self-sustaining actions, or their (organizational) identity, insofar it is the source of their individuality. This intrinsic logic, as opposed to some other, more common case, is a global form of coherence that is embodied by the system, in/by its permanently changing constitution, so that environmental selectivity (hence, interpretations and affordances) are not only specific to the system, but also historically dependent. Furthermore, while the notions of autopoiesis and autonomy have given us conceptual ways to think about life (and other phenomena by extension) in terms of self-persisting systems (Stano et al., 2023; Ward et al., 2017), further research, beginning from the more abstract notions of autonomy and organizational closure (Varela, 1979), in the context of an increasing shift in attention towards the material, biological nature of cognitive processes in Cognitive Science (van Gelder, 1998; Lyon, 2006; Engel et al., 2013), along with advances in Dynamical Systems Theory and computational implementations (Cliff et al., 1993; Beer, 1995, 1997; Harvey et al., 2005; Gershenson, 2007; Tani, 2017; Nolfi, 2021), contributed to a major boost in the artificial life agenda (Aguilar et al., 2014), which has made it possible to develop, relate and discuss theoretical properties such as primordial organismic traits (Ruiz-Mirazo and Moreno, 2004; Hanczyz and Ikegami, 2010; Froese et al., 2014), adaptivity (Di Paolo, 2005; Iizuka et al., 2013), sensorimotor habits (Buhrmann et al., 2013; Egbert and Barandiaran, 2014), normativity (Weber and Varela, 2002; Mojica, 2021) and precariousness (Beer and Di Paolo, 2023), among several others, all within a convergent (albeit not conflict-free (Barandiaran, 2017; Hutto and Myin, 2017)), more general framework for research in life and cognition.

Nonetheless, as is the case for most scientific endeavours, a better theoretical understanding often brings with it at least as many questions as the ones that it answers, be it through the formation of new concepts or the re-examination of others. The present work is an attempt at such a re-examinations – specifically, we propose the existence of a principled relation between organizationally closed and open-ended dynamics in adaptive autonomous systems.

In this respect, considering the many connotations associated with both these notions (Villalobos and Ward, 2015; Gershenson et al., 2020; Packard et al., 2019a,b), we shall briefly specify an interpretation, at least for the scope of our work, by making the following conceptual distinctions: First of all, as independently stated by Maturana (1987) and Varela (1984), it is important to make a clear-cut separation between the notion of self-organization in the rather cybernetic sense of a spontaneous aggregation of (physical, organismic or other) elements (Gershenson et al., 2020), and its misleading use associated with autopoiesis, which fundamentally denotes the preservation of an organizational identity. Likewise, we must distinguish autopoiesis, in its original sense of the defining property of living systems (Maturana and Varela, 1973), only possible in terms of (hence specific to) discrete molecular systems (Maturana, 2002), from broader notions of autonomy. Hence Varela (1979), in order to develop a concrete notion of autonomy beyond the cellular scope of autopoiesis (thus entailing the production of components and the need for a strict topological boundary), introduced the concept of organizational closure as a general property underlying the self-referential, circular nature of self-sustaining systems and, as such, as the fundamental trait defining the general category of autonomous systems (autopoietic systems being a specific minimal biological instance of these). In this sense, by assuming circularity as a core underlying principle, Varela (1979) reverses the order, thereby formally posing organizational closure as a necessary property for any autonomous system.

Along the same lines, it is vital to disentangle the more specific connotation of open-endedness that is conceptually tied to evolutionary processes, typically known as open-ended evolution (OEE) (Packard et al., 2019a,b), from open-endedness (OE) as a more general property (Stepney, 2021; Song, 2022). Whereas the former has become a prominent area of research in artificial life, in direct relation to other relevant phenomena (Ruiz-Mirazo et al., 2004; Witkowski and Ikegami, 2019; Borg et al., 2023) or essentially as a subject in its own right (Soros and Stanley, 2014; Soros, 2018; Pattee and Sayama, 2019; Corominas-Murtra et al., 2018), the latter, while certainly not neglected, has gathered less attention and it is still very much in a process of theoretical conceptualization (Stepney, 2021; Song, 2022). Furthermore, whilst we acknowledge the major role that open-endedness plays in evolution, we believe that insights about its basic principles prior to evolutionary considerations may be helpful for its integration into the overall artificial life research framework.

In this vein, we have decided to pursue our exploration as a minimal proof of concept making use of the Game of Life (GoL) cellular automaton (Gardner, 1970; Berlekamp et al., 1982). Like many before us, we consider GoL to be a well-suited toy scenario for an initial investigation, especially since, given the complexity involved, modeling biological processes in full detail is computationally unfeasible, but could also be theoretically counterproductive, insofar as, given our still very limited current knowledge, it would require too many assumptions and thus obscure any possible analysis. As a matter of fact, research through toy models has proven to be particularly fruitful for detailed investigation of several complex phenomena, as it permits conceptual demonstrations through characterizations of otherwise analytically intractable dynamics (Husbands, 2009; Beer, 2020c), and thus encompasses a rich literature (e.g., Varela et al. (1974); Wuensche (1994); Soros and Stanley (2014); Wang and Chan (2019); Hamon et al. (2022)). Particularly relevant to our work are the Bittorio model, introduced by Varela et al. (1991), and previous studies of gliders (and other transient patterns) in the Game of Life (Beer, 2004, 2014, 2015, 2020b,a) which have provided a relatable framework for investigation, mainly because observable patterns in the GoL emerge spontaneously from the dynamics of the cellular automata, while also displaying seemingly basic instances of properties such as self-organization, autonomy and so on.

Throughout the remainder of this paper we develop our main claim – that organizational closure and open-endedness act as complementary properties for any minimal autonomous system capable of some form of adaptation. For this purpose, we will initially present an elementary formal framework for exploring fundamental concepts related to organizational closure and autonomy, building towards its relation to open-endedness, we then illustrate these points through an experimental analysis of some of the transient dynamics of the Game of Life, to finally close with a brief discussion on possible implications, limitations and some concluding remarks.

A combined approach

In this section we will examine and attempt to clarify the fundamental properties of the theoretical notions introduced above by formalizing organizational dynamics in terms of multidimensional mappings, where state-to-state projections represent the structural transformations of the system co-determined by environmental factors (enactions). Building upon these ideas, we will relate the notions of organizational-closure and open-endedness in operational terms.

Organizational closure and closed domains

For organizational closure, unlike with material connotations, the domain is said to be closed because during its operation (structural instantiations), from a whole space of possible states, the system will only occupy a limited subset in which its viability conditions are not transgressed, transitioning from viable states into (new or not) viable states, thus defining an operational subdomain. And so, the organizationally closed domain becomes the domain of existence as given by the intrinsic autonomous logic of the system.

Organizationally closed systems, in this sense, are systems that, because of their particular recursive nature, will always yield state transitions leading to states that are a structural instance within the set of states of the same system and that, in turn, will therefore lead to further states with this same property and so on and so forth. This is akin to the mathematical sense of closure, whereby performing an operation on any element of a (sub)set will result in some

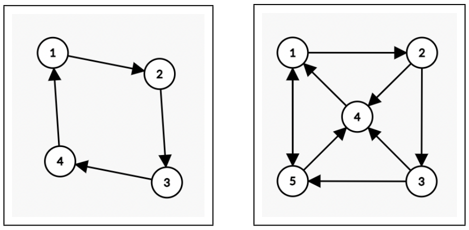

Figure 1: To the left, a minimal example diagram of an operationally closed system. Nodes represent states; arrows, the transitions. Circularity (i.e. closure) is given by the recursion of the state transitions. To the right, a diagram of another, a bit more complex operationally closed system. In this case, closure is also a product of the recursive transitional dynamic, although it is not merely linear, because of the different responses available to each state.

element of the same (sub)set; where the set is the (potentially unlimited) set of states and the operation is the state-transition. These ideas are sketched in figure 1 to underline the notion of closure through recursion, and to illustrate the effect of different transitional dynamics on a system, entailing different responses (transitions) according to the particular (structural and environmental) circumstances in which this process takes place.

The system coherent state transitions can also be interpreted as (cognitive, although with a purely mechanistic connotation in our context) distinctions, or as a form of intelligence, insofar as each structural change is considered a system-environment correlated response. In other words, any organizationally closed system could be conceived to be intelligent at least to the extent that it is capable of, by means of its structural coherence, systematically determining a correct (viable) successive state from an otherwise indistinct flux of external changes. This primary form of intelligence, which in the enactive literature is fundamentally understood as embodied (i.e., there are no abstract symbolic instructions, just the structure of the system, or the body of an organism), could be equated to – or expressed as – the system's selectivity. Nevertheless, selectivity is only useful as a concept, inasmuch as the system upon which it's operating is organizationally closed. That is to say that any distinction (structural transformation) by means of selectivity is logically pertinent if and only if there is a system with an intrinsic input, defined by its global recursive dynamics, to which the difference of making some distinction instead of another has any relevance at all. Along these lines, an enaction can be defined as a manifestation of these structurally encoded material constraints stemming from a broader system-environment organizational coherence that steers the continuous transformation of the system (its behavior) towards self-sustaining states. And it is in this sense that we believe it is the core element upon which mappings characterizing the organizational closure of a system should

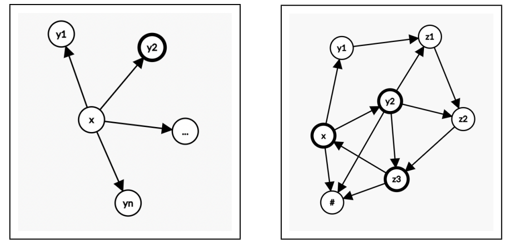

Figure 2: To the left, a minimal example diagram of the selectivity of a system. It is considered to be a primary manifestation of intelligence or of a cognitive distinction, because solely on the basis of its structural properties, an appropriate response to environmental circumstances is given. To the right, a sample organization of a system. The imaginary sequence undergone by the hypothetical system is $x \rightarrow y2 \rightarrow z3 \rightarrow x$ (thicker nodes), the hash node represent the destruction of the system, potentially following from any visited state. Every arrow represents a transition connecting two states and their conjunction an enaction.

be built (see figure 2 for a minimal depiction of the notion of selectivity and a sketch of its role in the organized system dynamics).

A final comment before proceeding onto the next section: While we have referred to the notion of valid or invalid transitions in the deterministic sense (given the context of cellular automata), this should not mistakenly imply the idea of binary outcomes (the latter resulting in the direct disintegration of the system) for other kinds of systems.

Open-ended closure?

Although seemingly contradictory from a first impression, there are two considerations that need to be taken into account; first, previously unknown environmental circumstances may produce transitions which, albeit valid, may direct the structural transitions towards non previously instantiated configurations. Second, given that structural selectivity is specific and dynamic, the aggregation of new valid organizational nodes may determine new behavioral pathways. In simple words, whenever faced with the unknown, either the system will disintegrate or produce an original form of recursion, be it as a new structural instance of itself, as new state transition, or both. If successful, this enlargement of the organizational domain not only will increase robustness/resilience through structural discrimination, but more importantly, it may be able to provide a whole new series of dynamical transformations that can even surpass the dynamical pull from the parent domain. Moreover, since the number of transitions among viable states is far from uniform, the general organizational domain will give place to more and less common behavioral trajectories, which can be topographically represented as dynamical attractors (these ideas are very simply depicted in 3).

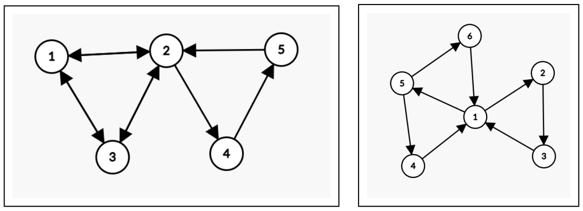

Figure 3: To the left: the original organizational space made of states 1, 2 and 3 creates a new node (state 4). From this, another node is created (state 5) giving way to a larger organizational space. Future environmental contingencies could result in further expansions stemming from either states 4 or 5. To the right: given the organizational dynamics, many more transitions will lead to state 1 than to other states. Hence, dynamically speaking, state 1 can be said to act as an attractor within the organizational closed space.

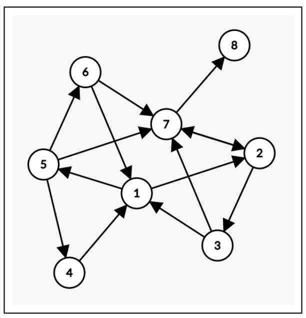

Figure 4: The creation of a new node may produce new transitions and hence, a topological redistribution of the organizational space.

The effect of new nodes or/and projections is then the rearrangement of such dynamics in ways that could be quite unpredictable. A trivial example of a hypothetical case is presented in figure 4.

From our perspective, these two properties are equally necessary for any minimal form of autonomy. Fundamentally, open-endedness requires something to remain in order to manifest itself. In very abstract terms this could be understood as an entity, a general property or, specifically, as we are posing here, an organizationally closed domain. Otherwise it wouldn't be much more than evanescent fluctuations without any long-term causal implication. Likewise, a static or fixed organizational domain would be devoid of any capacity for adaptation, and even after a fortunate set of coincidences from which some organizationally recursive system could have spawned, even the slightest change in the environmental circumstances would trigger its demise. Basically then, our proposal is that open-ended organizationally closed systems are a primary kind of autonomous systems.

By open-ended we don't intend a necessarily beneficial or virtually infinite connotation somehow akin to the idea of a classical search algorithm, but rather the opposite; a quite structurally and organizationally constrained property which is intrinsic, unavoidable and that does not have any purpose in itself, in the exact same way as the recursive dynamics that enable the organizational closure of a system. As a matter of fact, from our point of view, it is precisely because of this that the notion of an open-ended closure can be conceptually tied to autonomy; because it provides a more naturalized, plainly mechanistic account of minimal adaptive capacities that are non optional, that unfold historically and coherently to the extent that the structural modulation of the system can withstand it and that do not involve any form of semantic or incipient proto-mental properties underpinning its operation.

In the following section we will try to illustrate these concepts through a toy case involving experimental analysis built around the dynamics of transient patterns present in Game of Life cellular automata.

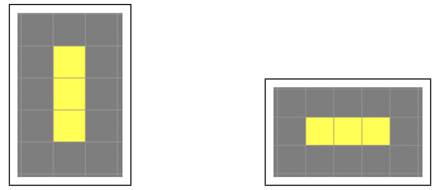

Figure 5: The two instances of the dynamic self-sustaining activation pattern known as a Blinker. These two forms are recursively produced by each other indefinitely in the absence of environmental perturbations.

Open-ended organizational closure in the Game of Life

Following on from previous studies from Beer (2004, 2014, 2015, 2020b,a) we will consider any given identifiable pattern to be the combined active cells plus the non-active cells in its Moore neighborhood (the surrounding cells that contribute to the update of the cell values at every time-step), which can be conceived of as a functional membrane. Along the same lines, we will consider valid organizational transitions those processes whereby an identifiable pattern updates its structure as a whole, either into the same, or into a different (valid) structural instantiation. Accordingly, invalid transitions will be any causing the system to disintegrate.

The minimal case: a Blinker

The GoL pattern known as a blinker (fig. 5) is the minimal self-sustaining dynamical pattern that can be encountered in the Game of Life. As such, it is probably a good starting point for our analysis.

Since the blinker can only display two structural instances, its organizational space in an empty grid could be represented as two states recursively producing each other. However, when taking into account the fact that both these structural instances are equivalent under a simple rotation

(Beer, 2004) this can ultimately be reduced to a single self-connected state. Whereas in an empty grid there is no possible disintegration, when accounting for all the possible causal effects of the environment over the blinker, we calculate that from the whole range of possible combinations of environmental cell states $2^{20} = 1,048,576$, only $166,638$ will reproduce a blinker (either horizontally or vertically oriented), hence making its hypothetical chances of being sustained around 6%. Thus we could express the blinker's valid transitional dynamics as:

$$(B_x, E_x) \longrightarrow B_y \quad \text{where: } B_x = B_y$$

Where B_x and B_y stand for the valid instances of the blinker, while E_x denotes the environmental category set that contains all the grid environmental configurations enabling this transition.

We shall at this point, however, consider an issue that is central to our thesis and that might be evident by now, namely: how can we be sure that the only structural instance that belongs to the organization of the blinker is the one we just described? First of all, what we denominate a blinker is something that is extrinsically defined, namely, a group of 3 active cells forming an active region shaped as a row (plus the surrounding inactive cells acting as a membrane). However, in spite of being intuitively clear, this distinction is to a great extent a perceptual impression or somehow an illusion that is strongly biased, and it is underpinned by a premise of structural similarity that may be misleading. This is somehow like identifying patterns and shapes when looking at the stars or at the clouds in the sky; it is useful for observation and to begin a categorization, but it may prove deceitful as a systematizing principle. To be clear, the point that we would like to stress is that, apart from a rather problematic identification based on resemblance between patterns, the fact is that the only formal criteria that we have, at least for now, to ascertain that some given pattern belongs to an organizationally closed domain, is to check if it display a recursive transitional behavior leading it back to some of the known structural configurations after a given number of time-steps, in other words: to examine if its aggregation is coherent with the recursive organizational dynamics of the system unfolding over time.

Simply put, our view is that, if hypothetically speaking, we were able to find a sequence of activation patterns such that, starting from a canonical blinker (or from any other transient pattern for that matter), it eventually leads back to the blinker (hence, another instance of a cyclical attractor in its own organizational space), then we should conceive the activation patterns of this sequence as variants, even if less intuitive, of the canonical blinker case, therefore also as belonging to the set of structural instances of a more general organization. Put another way:

$$(B_x, E_x) \longrightarrow (U_1, e_1) \longrightarrow ... \longrightarrow (U_n, e_n) \longrightarrow (B_y, e_y)$$

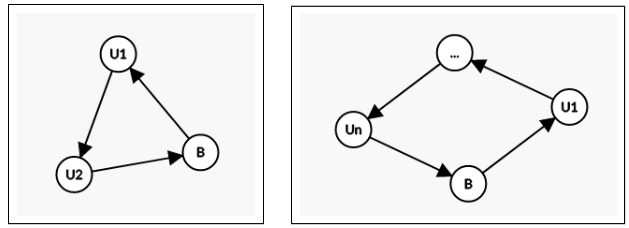

Figure 6: Graph representations of different possible aggregation of structures as recursive transitions starting from and ending in a canonical state

Where by $U_1, U_2, ..., U_n$ we refer to unknown structural cases that enable an organizational progression, extending the recursive domain. This can also be represented more graphically as in figure 6, where a structural aggregation into the recursive organizational transitions are presented.

In this case, by expanding our notion of what the organization of a blinker may be, we would move from a recurrent self-directed dynamic with only one valid structural instance, onto a wider mapping allowing for a set of structural instances that participate in multiple enactions. Thus, by taking into account other environmental circumstances, we wouldn't only have a more complex transitional mapping from- and to- the blinker, but also a more complete description of what the main system and its general organization really is.

A multiplicity of organized structures

In order to test this idea of an expanded domain, we first simulated all the possible $2^{20} = 1,048,576$ environmental cases for the blinker and analyzed the resulting domains, after transitions. However, given the possible combinatorial explosion and since our goal is rather a proof of concept instead of a full characterization, we constrained our search for patterns in the co-domain, only to states having between 3 and 5 active cells (initially; 68,315 cases). As we will see, this turns out to be enough for our theoretical exploration.

Now then, given that this is still quite a high volume of instances to examine, we applied two methods to systematically analyze them: first, to avoid domain/structure confusions, we ran two filters as a pre-processing step; a domain-filter, whereby we removed all isolated active cells (i.e., those that weren't part of any pattern themselves) and a decay-filter, where we excluded all the unstable domains from the remaining set (those that will unavoidably disintegrate in the next time-step). This left us with a total of 129,205 cases. Subsequently, taking advantage of the fact that, as it has been proposed and characterized in Beer (2004, 2014), active configurations in the Game of Life can be reduced to equivalent structures by applying different types of symmetries, namely: rotation, translation and transposition, we clustered the resulting patterns into equivalent types under these operations. From this we obtained a total

Active cells	Valid transitions	Domain filter	Decay filter	Equivalent types
3	5,901	55,541	15,175	2
4	20,174	57,228	53,410	15
5	42,240	60,838	60,620	54
total	68,315	173,607	129,205	71

Table 1: Number of transitional patterns obtained from a blinker embedded in different environmental circumstances. Results are classified according to the number of active cells. Domain and Decay filters actually increase the number of transitions to examine, by removing active cells unrelated to the structural patterns. Nevertheless, there is drastic reduction after accounting for symmetries, as shown in the last column (further detail in text)

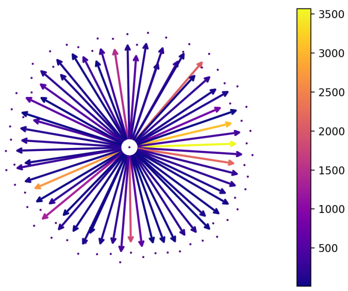

Figure 7: Elementary graph representation of the valid potential transitions that a blinker can perform without disintegrating (apart from the transition into itself). Each node represents a symset (set of equivalent structures under symmetry operations) and colors represent the 'weight' or the number of environments that produce a given transition.

of 71 symmetrical sets, or symsets. (see table 1 and fig.7 for complementary information).

While it may be interesting to know how many non-destructive transition the environment can trigger (co-determine), these will only be relevant to the extent to which they provide a supporting organizational pathway; they must eventually lead to some other valid structural node. Thus, we applied the same filtering and reduction process, iterating over a subset of the patterns resulting from the blinker, looking for possible cases of recursion (see fig.8 for a list of these patterns). The results are presented in table 2.

Conceptually speaking, every symset represents a transformation into a viable state whereby the structural selectivity of the system responds to a variety of environmental circumstances as if they were the same and which, at least functionally, work as such. Since each of these distinctions is equally motivated by the range of alternatives that the structural instantiation of the system is capable of encoding, as well as by the variability of the environment, more resilient systems should exhibit correct state transitions encompassing a larger environmental subset. This, however, is not so straightforward as it appears because of the unavoidable trade-off between closure and openness. While the number of environments that a system can handle is hard to question as an objective factor for sustainability, the man-

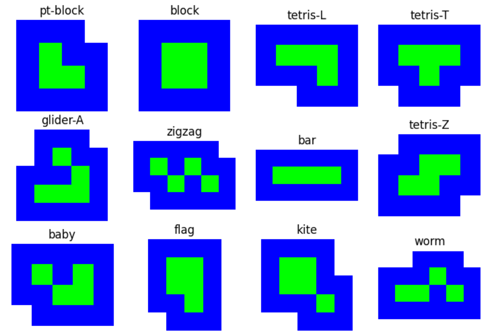

Figure 8: Twelve of the fourteen GoL patterns (apart from the blinker and the second structural instance of the glider) examined in the experiments described in this section. In order to perform a broader analysis of possible organizational transformations, each one of them was embedded into the whole set of environmental configurations for transitions to be simulated and later reduced through filtering and symmetry operations (see table 2).

Structural Pattern	AC 3	AC 4	AC 5	Equivalent Types	Valid $x \to y$
blinker	2	15	54	71	129,205
pt-block	1	13	34	48	10,421
block	1	8	23	32	4,867
tetris-L	2	15	64	81	73,124
tetris-T	0	1	6	7	1,704
glider-A	1	19	74	95	186,852
glider-B	0	10	54	64	27,805
zigzag	0	1	2	3	75,765
bar	2	14	40	56	324,518
tetris-Z	0	1	6	7	1039
baby	2	19	65	86	64,342
flag	2	23	86	111	63,134
kite	2	19	74	95	30,243
worm	1	6	27	34	108,157

Table 2: Number of (valid) equivalent types (after symmetry operations) and corresponding valid transitions instances for each structural pattern analyzed, according to their number of active cells (AC) (Further detail in main text).

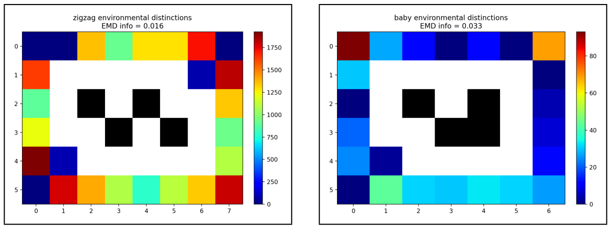

Figure 9: Visual representation of GoL grid subdomains in terms of environment-structure specificity. The colored environmental cells depict the (probability) distribution from the equivalent set of all the environmental configurations that will produce a same structural response. Dark colours represent fewer active instances; conversely, light colours, a high probability of activation for that cell. As can be seen, the zigzag (left) interpretation is far less precise (i.e., more homogeneous, hence more uncertain) than that of the baby pattern (right) where two cells (left/right top corners) can be interpreted as cues. This can be measured as information.

ner in which these responses are distributed, in terms of innovation versus control, is also highly relevant, especially when taking into account the limited algorithmic capacity that GoL patterns can embody (this probably extends to simple organisms, in cognitive terms). A clear example of this is the opposite symsets/environments relations displayed by the 'zigzag' and 'baby' patterns (see table 2). Whereas the former has only three different behavioral alternatives (structural transitions), it still accounts for a large range of external cases; conversely, 'baby' (likewise 'flag' or 'tetris-T') possesses many more available (structural) actions, for an overall similar number of circumstances to which it can apply them. This can be represented visually and in terms of information, by contrasting the probability distribution associated with the set of possible environmental states that produce the same specific transitions. For this, we specifically used the Earth Mover's Distance (EMD) (Rubner et al., 1998) (see example in fig.9).

Our guess is that organizational dynamics modulate these tendencies in a general fashion, rather than moving towards a local (by structural instantiation) equilibrium. As we know, behavioral convergence has the great advantage of exerting some degree of control upon external dynamics, making them more predictable (canonical patterns in empty grids serve as a good analogy for this: as long as the canonical cycle continues, canonical structures recursively produce others, but more importantly, they don't mess-with/alter the environment emptiness). However, assuming an axiomatic higher environmental potential variability than a system contained by it, distribution into more branches can provide the system with more chances for sustaining than even a hypothetically full, but narrow environmental categorization. It is in this sense that we claim that organizationally closed (or autonomous) systems continuously and progressively ex-

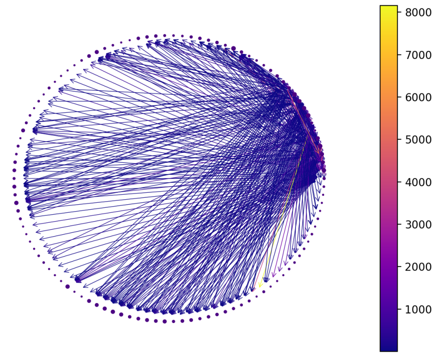

Figure 10: The full mapping of valid transitions spawning from the patterns discussed in this section. While open-ended nodes (i.e., of uncertain/unknown further transformations) account for most cases, there is a clear recursive component giving rise to an expanded organizational domain.

pand their organizational domains in an open-ended fashion, because system-environment interactions impose transformations that, if not disintegrative, will produce new forms of structural coherence, therefore altering the intrinsic organizational logic. Since these extrinsic effects can't be avoided, resilient systems would then be those capable of self-sustaining by a robust integration of aggregated states into the organizational domain through production of further recursive transitions. Exactly this kind of process seems to take shape in the organizational space of the expanded blinker (see fig.10 for reference), where open (previously unknown) states from the full valid mapping of blinker transformations transition into other unknown states (as well as known ones) under given external constraints.

The instantiation of recursive structural relations can be appreciated better by considering just the transitions between the patterns that have been exhaustively analyzed (within the cap for number of active cells), this is displayed in figure 11. From this more general stance, it is difficult not to consider this (limited) set of structures as if they were not organizationally overlapping, insofar as transitions among states are both selectively specified and environmentally codependant.

Discussion and Concluding remarks

Throughout this paper we have presented an approach to the notion of organizational closure in terms of multidimensional mappings, whereby the closed domain produced by the structural transformations of the system can be characterized and, building upon this, we have conceptually related organizational closure with the (general) notion of open-endedness in terms of adaptive autonomous dynamics. Finally, we have illustrated these ideas within a toy experiment setup in GoL, re-examining previous assumptions regarding structural and organizational correspondences (Beer, 2004,

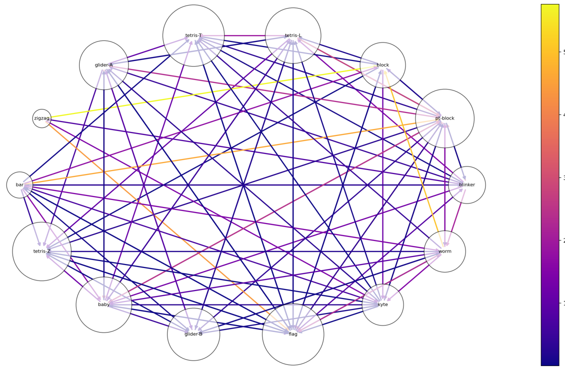

Figure 11: Mapping of the transitions among the 12 patterns examined through this section. Considering that these transitions obey the differential structural selectivity given by the concrete configuration of the patterns, as wells as the environmental circumstances in which the structure is embedded, it seems reasonable to conceive them as different instances belonging to the same organization.

2014, 2015). For reasons of space and computational limitations, it is infeasible to provide an exhaustive characterization, nonetheless, the exploration of the transitional dynamics of the set of GoL patterns presented here consistently show how open-ended organizational aggregation is extrinsically imposed due to unpredictable environmental fluctuations that can co-specify novel state transitions, expanding the organizational domain of the system and its otherwise extremely precarious self-sustaining dynamics. Insofar as structural configurations that are previously unknown to the system are produced, this can be interpreted as a form of (non-evolutionary) innovation. Moreover, while the notion of innovation may be conceptually obscure (Stepney, 2021), by framing it in terms of an organizational space it becomes operationally distinguishable; concretely speaking, rather than the re-instantiated structure in itself (i.e., considered in organizational isolation), the crux is whether the organizational space is reshaped by it, be it by a new transformation, or by a new found transition between known structures.

Two main limitations we should mention are the use of structures with no more than five active cells and the synchronous update imposed by GoL rules. While the former may be solved with enough computational resources (although our guess is that the main principles described here would hold), the latter requires further analytical and conceptual considerations, in order to fully incorporate asynchronous dynamics and multiple interacting time scales. In fact, we believe that this could be a productive line for future work, as the role of this kind of temporal dynamics has been theoretically related to more complex cognitive properties (Varela, 1999; Dorato and Wittman, 2015; Friston, 2018; Rodriguez et al., 2023).

An interesting, probably inevitable question is whether the whole set of possible patterns of activation could be conceived as just one organizationally closed system. We reckon that this will depend to a large extent on how we define identity among instances. While we have considered structural recursion (i.e., the re-instantiation of a given structure or a sequence of them) as a criterion for closure (given our specific toy experiment and the nature of the GoL), it is important to keep in mind that organizationally closed dynamics don't necessarily require this to be the case, something which is evident when considering the development of living organisms, in which these structural nodes or attractors are replaced by functions upon viability conditions in a more complex state space. An alternative criterion could be individualization through membranes, which has also been proposed as a necessary condition (in more or less abstract ways, e.g. Maturana and Varela (1973); Varela (1979); Beer (2004); Di Paolo et al. (2017)). Although this is to a large extent implicit in our search method, a comprehensive exploration starting from this premise would include other cases (such as nested patterns, for example) and in consequence, a different methodology.

Regarding the elusive relation between autonomous dynamics and adaptive capacities, while the notion of adaptivity has gained relevance as an alternative to the rather adaptive-less original formulation of autopoiesis (Di Paolo, 2005; Di Paolo et al., 2017), conceptual nuances entailed by its reliance on semantic attributions cast doubts on it (Hutto and Myin, 2017), at least in the context of minimal autonomous systems. This is a second line we envision for future work and we hope that the notion of open-ended organizational closure proposed here may contribute as a sort of theoretical middle ground. Whether we conceptualize it as a single property (insofar as closure and openness only manifest consonantly) or as two properties in a continuous oscillatory tension, in future work it may become important when analyzing evolutionary scenarios, especially with respect to the transmission of genetic information (Soros and Stanley, 2014; Packard et al., 2019a,b). In this sense, regarding possible logical orders of implication (Hintze, 2019; Pattee and Sayama, 2019), we guess that, like in many natural cases, there is a constant trade-off between stability and flexibility. The former being fundamental to self-sustained existence and probably incremental in the absence of significant environmental changes, whereas the latter provides the system with the potential to overcome these changes through the rearrangement of its organizational dynamics. And although advantages and disadvantages of both these properties are relatively evident, in general there won't be something that we can point at as some sort of natural tendency or preference, given that the modulation between them will actually fluctuate continuously depending mainly on the environmental context, rather than on some intrinsic disposition of the system towards innovation as a purpose in itself.

References

Aguilar, W., Santamaria-Bonfil, G., Froese, T., and Gershenson, C. (2014). The past, present and future of artificial life. *Frontiers in robotics and AI*, 1(8).

Barandiaran, X. (2017). Autonomy and enactivism: Towards a theory of sensorimotor autonomous agency. *Topoi*, (36):409–430.

Beer, R. (1995). A dynamical systems perspective on agent-environment interaction. *Artificial Intelligence*, 72:173–215.

Beer, R. (1997). The dynamics of adaptive behavior: A research program. *Artificial Intelligence*, 20:257–289.

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

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

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

Beer, R. (2020a). Bittorio revisited: Structural coupling in the game of life. *Adaptive Behavior*, 28(4):197–212.

Beer, R. (2020b). An integrated perspective on the constitutive and interactive dimensions of autonomy. In *Proceedings of the ALIFE 2020: The 2020 Conference on Artificial Life*, pages 202–209.

Beer, R. (2020c). Lost in words. *Adaptive Behavior*, 28(1):19–21.

Beer, R. and Di Paolo, E. (2023). The theoretical foundations of enaction: Precariousness. *Biosystems*, 223(104823).

Berlekamp, E., Conway, J., and Guy, R. (1982). *Winning ways for your mathematical plays, vol. 2*. New York: Academic Press.

Borg, J., Buskell, A., Kapitany, R., Powers, S., Reindl, E., and Tennie, C. (2023). Evolved open-endedness in cultural evolution: A new dimension in open-ended evolution research. *Artificial Life*, pages 1–22.

Buhrmann, T., Di Paolo, E., and Barandiaran, X. (2013). A dynamical systems account of sensorimotor contingencies. *Frontiers in Psychology*, 4(285):1–19.

Cliff, D., Husbands, P., and Harvey, I. (1993). Explorations in evolutionary robotics. *Adaptive Behavior*, (2):73–110.

Corominas-Murtra, B., Seoane, L., and Solé, R. (2018). Zipf's law, unbounded complexity and open-ended evolution. *J. R. Soc. Interface*, 15(149).

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

Di Paolo, E., Burhmann, T., and Barandarian, X. (2017). *Sensorimotor Life: An enactive proposal*. Oxford University Press.

Dorato, M. and Wittman, M. (2015). The now and the passage of time. *KronoScope*, 15:191–213.

Egbert, M. and Barandiaran, X. (2014). Modeling habits as self-sustaining patterns of sensorimotor behavior. *Frontiers in Human Neuroscience*, 8(590):1–15.

Engel, A., Maye, A., Kurthen, M., and König, P. (2013). Where's the action? the pragmatic turn in cognitive science. *Trends in Cognitive Science*, 17(5).

Friston, K. (2018). Am i self-conscious? (or does self-organization entail self-consciousness?). *Frontiers in Psychology*, 9.

Froese, T. and Di Paolo, E. (2011). The enactive approach. theoretical sketches from cell to society. *Pragmatics and Cognition*, 19(1):21–36.

Froese, T., Virgo, N., and Ikegami, T. (2014). Motility at the origin of life: Its characterization and a model. *Artificial Life*, 20(1):55–76.

Gardner, M. (1970). Mathematical games: The fantastic combinations of john conway's new solitaire game 'life'. *Scientific American*, 223:120–123.

Gershenson, C. (2007). *Design and Control of Self-Organizing Systems*. Mexico City: CopIt Arxives.

Gershenson, C., Trianni, V., Werfel, J., and Sayama, H. (2020). Self-organization and artificial life. *Artificial Life*, 26:391–408.

Hamon, G., Etcheverry, M., Chan, B., Moulin-Frier, C., and Oudeyer, P. (2022). Learning sensorimotor agency in cellular automata. *https://developmentalsystems.org/sensorimotor-lenia*.

Hanczyz, M. and Ikegami, T. (2010). Chemical basis for minimal cognition. *Artificial Life*, 16(3):233–243.

Harvey, I., Di Paolo, E., Wood, R., and Quinn, M. (2005). Evolutionary robotics: A new scientific tool for studying cognition. *Artificial Life*, 11:79–98.

Hintze, A. (2019). Open-endedness for the sake of open-endedness. *Artificial Life*, 25(2):198–206.

Husbands, P. (2009). Never mind the iguana, what about the tortoise? models in adaptive behavior. *Adaptive Behavior*, 17(4):320–324.

Hutto, D. and Myin, E. (2017). *Evolving Enactivism. Basic Minds Meet Content*. MIT Press.

Iizuka, H., Ando, H., and Maeda, T. (2013). Extendeed homeostatic adaptation model with metabolic causation in plasticity mechanism-toward constructing a dynamic neural network model for mental imagery. *Adaptavie Behavior*, 21(4):263–273.

Lyon, P. (2006). The biogenic approach to cognition. *Cognitive Processing*, 7(1):11–29.

Maturana, H. (1987). Everything said is said by an observer. In *Thompson W. I. (ed.) Gaia: A way of knowing*. Lindisfarne Press, New York.

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

Maturana, H. and Varela, F. (1973). *Autopoiesis: the organization of the living. [De maquinas y seres vivos. Autopoiesis: la organizacion de lo vivo]. 7th edition from 1994.* Editorial Universitaria.

Mojica, L. (2021). The enactive naturalization of normativity: from self-maintenance to situated interactions. *History and Philosophy of Life Sciences*, 43(127).

Nolfi, S. (2021). *Behavioral and Cognitive Robotics: An Adaptive Perspective.* Roma, Italy: Institute of Cognitive Sciences and Technologies, National Research Council (CNR-ISTC).

Packard, N., Bedau, M., Channon, A., Ikegami, T., Rasmussen, S., Stanley, K., and Taylor, T. (2019a). Open-ended evolution and open-endedness. editorial introduction to the open-ended evolution i special issue. *Artificial Life*, 25(1):1–3.

Packard, N., Bedau, M., Channon, A., Ikegami, T., Rasmussen, S., Stanley, K., and Taylor, T. (2019b). Open-ended evolution and open-endedness. editorial introduction to the open-ended evolution ii special issue. *Artificial Life*, 25(2):93–103.

Pattee, H. and Sayama, H. (2019). Evolved open-endedness, not open-ended evolution. *Artificial Life*, 25(1):4–8.

Rodriguez, F., Husbands, P., Ghosh, A., and White, B. (2023). Frame by frame? a contrasting research framework for time experience. In *ALIFE 2023: Ghost in the Machine: Proceedings of the 2023 Artificial Life Conference*, page 75.

Rubner, Y., Tomasi, C., and Guibas, J. (1998). A metric for distributions with applications to image databases. In *Proceedings of the 1998 IEEE International Conference on Computer Vision, Bombay, India*.

Ruiz-Mirazo, K. and Moreno, A. (2004). Basic autonomy as a fundamental step in the synthesis of life. *Artificial Life*, 10(3):235–259.

Ruiz-Mirazo, K., Peretó, J., and Moreno, A. (2004). A universal definition of life: Autonomy and open-ended evolution. *Origins of Life and Evolution of the Biosphere*, (34):323–346.

Song, A. (2022). A little taxonomy of open-endedness. In *Workshop on Agent Learning in Open-Endedness (ALOE) at ICLR 2022*.

Soros, L. (2018). *Necessary Conditions for Open-Ended Evolution*. PhD thesis, Electronic Theses and Dissertations, 2004-2019. 5965.

Soros, L. and Stanley, K. (2014). Identifying necessary conditions for open-ended evolution through the artificial life of chromaria. In *ALIFE 14: Proceedings of the Fourteenth International Conference on the Synthesis and Simulation of Living Systems*, pages 793–800.

Stano, P., Nehaniv, C., Ikegami, T., Damiano, L., and Witkowski, O. (2023). Autopoiesis: Foundations of life, cognition and emergence of self/other. *Biosystems*, 232(105008).

Stepney, S. (2021). Modelling and measuring open-endedness. In *OEE4 workshop, at ALife 2021, Prague, Czech Republic (online), July 2021*.

Tani, J. (2017). *Exploring Robotics Minds. Actions, Symbols, and Consciousness as Self-Organizing Dynamic Phenomena*. Oxford University Press.

van Gelder, T. (1998). The dynamical hypothesis in cognitive science. *Behavioral and Brain Sciences*, 21:615–665.

Varela, F. (1979). *Principles of Biological Autonomy*. North Holland.

Varela, F. (1984). Two principles for self-organization. In *Ulrich, H., Probst, G.J.B. (eds.) Self-Organization and Management of Social Systems*. Springer Series on Synergetics, vol 26. Springer, Berlin, Heidelberg.

Varela, F. (1999). The specious present: A neurophenomenology of time consciousness. In *J. Petitot, F. J. Varela, B. Pachoud, & J.-M. Roy (Eds.), Naturalizing phenomenology: Issues in contemporary phenomenology and cognitive science*. Stanford University Press.

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

Varela, F., Thompson, E., and Rosch, E. (1991). *The embodied mind: Cognitive science and human experience*. The MIT Press.

Villalobos, M. and Ward, D. (2015). Living systems: Autonomy, autopoiesis and enaction. *Philosophy & Technology*, 28(2):225–239.

Wang, B. and Chan, C. (2019). Lenia - biology of artificial life. *arXiv*, (182.05433v3).

Ward, M., Silverman, D., and Villalobos, M. (2017). Introduction: The varieties of enactivism. *Topoi*, 36:365–375.

Weber, A. and Varela, F. (2002). Life after kant: Natural purposes and the autopoietic foundations of biological individuality. *Phenomenology and the Cognitive Sciences*, 1:97–125.

Witkowski, O. and Ikegami, T. (2019). How to make swarms open-ended? evolving collective intelligence through a constricted exploration of adjacent possibles. *Artificial Life*, 25(2):178–197.

Wuensche, A. (1994). The ghost in the machine: Basins of attraction of random boolean networks. In *Artificial Life III: SFI Studies in the Sciences of Complexity vol.XVI*.