The Cognitive Domain of a Glider in the Game of Life

Randall D. Beer*
Indiana University

Abstract   This article examines in some technical detail the application of Maturana and Varela's biology of cognition to a simple concrete model: a glider in the game of Life cellular automaton. By adopting an autopoietic perspective on a glider, the set of possible perturbations to it can be divided into destructive and nondestructive subsets. From a glider's reaction to each nondestructive perturbation, its cognitive domain is then mapped. In addition, the structure of a glider's possible knowledge of its immediate environment, and the way in which that knowledge is grounded in its constitution, are fully described. The notion of structural coupling is then explored by characterizing the paths of mutual perturbation that a glider and its environment can undergo. Finally, a simple example of a communicative interaction between two gliders is given. The article concludes with a discussion of the potential implications of this analysis for the enactive approach to cognition.

1   Introduction

Since the publication of its original manifesto over 20 years ago [28], interest in what has come to be known as the *enactive approach* to cognition has been steadily growing [10, 22, 24]. This approach is grounded—historically, conceptually, and even terminologically—in earlier work by Maturana and Varela (henceforth MV) on the biology of cognition [15, 16, 25]. Unfortunately, MV's framework can be difficult to understand, requiring a fundamental shift in perspective that many find strange on first reading. Even among experts, there are a number of disagreements of interpretation that have not yet been fully resolved. Thus, it would seem that a detailed analysis of a concrete example of MV's framework in action could be very beneficial to ongoing work in enactive cognitive science. Such is the goal of this article.

There is a long history of developing computational models of the concept of autopoiesis that serves as the foundation of MV's biology of cognition [3, 9, 17–19, 27, 29, 33]. There have even been attempts to engineer physical autopoiesis [13]. However, the focus of most of this work has been on simply demonstrating that autopoiesis can be produced; there has been very little systematic mathematical analysis of such models. More importantly for our purposes here, most work on modeling autopoiesis has not directly engaged the behavioral and cognitive implications of MV's biology of cognition, although there are important recent exceptions [6–8, 23]. Building on a previous proposal [1], the central aim of this article is to rigorously and exhaustively characterize the

---

* Cognitive Science Program, 1900 E. 10th Street, 840 Eigenmann, Indiana University, Bloomington, IN 47406. E-mail: rdbeer@indiana.edu

complete cognitive domain of a glider in the game of Life, as well as the structure of the interactions in which it can participate with its environment.

This article is organized as follows. Section 2 introduces the key concepts and terminology of MV's framework that we will be exploring. Section 3 then reviews the game of Life cellular automaton and argues for the utility of considering gliders from an autopoietic perspective. Section 4 characterizes the cognitive domain of a glider by enumerating its reaction to all possible perturbations that it can receive from its immediate environment. In Section 5, we examine the structure of the different response classes that result from these perturbations and how this structure is grounded in a glider's constitution. Sections 6 and 7 then begin a characterization of the ways in which a glider and its environment can interact over time, while Section 8 presents a preliminary investigation of communicative interactions between gliders. Finally, we conclude in Section 9 with a discussion of some of the broader issues raised by this analysis and suggest some directions for future work.

2   Autopoiesis and Cognition

Let us begin with a brief introduction to the key concepts of MV's framework that we will need in this article. This section makes no claim to being exhaustive and, in the interests of clarity, many subtleties are purposely simplified or ignored. For more complete presentations, see the original literature [15, 16, 25] or more recent discussions [24].

With one crucial exception, MV's framework derives from a fairly traditional perspective on science, albeit one expressed in a somewhat idiosyncratic language. Scientists can only study systems that can be distinguished in some way from the general background of experience; MV call such a system a *simple unity*. For example, a sealed pressure vessel containing a gaseous mixture is a simple unity. Simple unities have properties (e.g., volume, temperature, pressure, and color) that can be measured and manipulated, as well as relations between these properties (e.g., the ideal gas law and Planck's law of blackbody radiation). But scientists can also treat such a system as a *composite unity* by further decomposing it into a collection of components (e.g., the molecular constituents of the gas and vessel) with their own properties and relations. Scientific explanation typically involves building bridges between the properties and relations of a unity and the properties and relations of its components (e.g., the derivation via statistical mechanics of the ideal gas law and the quantum mechanical derivation of Planck's law).

MV make an important distinction between the *organization* of a composite unity and its *structure*. A composite unity's organization is given by the abstract relations between its components that are essential to its being the particular kind of unity that it is. A composite unity's structure, on the other hand, is given by the actual components and relations between them that constitute a concrete instance of that unity. For example, for a system to be a pressure vessel, the molecules of the vessel must satisfy relations of integrity and containment with respect to the molecules of the gas; these relations define its organization. In contrast, the additional component properties and relations necessitated by any concrete realization of a pressure vessel constitute its structure. For example, the types, positions, orientations, and velocities of the molecular constituents of an actual pressure vessel are part of its structure. Note that MV argue that any scientifically explainable composite unity must be a *structure-determined system*; its behavior must be determined ultimately by the complete set of concrete properties and relations of its particular instantiation.

The structure of a composite unity is always changing, both as a result of its own intrinsic dynamics and as a result of any external perturbations that are applied to it. For example, the types, positions, orientations, and velocities of the gas molecules in a pressure vessel are in constant flux as they career about the container, colliding with one another, ricocheting off the walls of the vessel, and making and breaking chemical bonds. In addition, we can externally perturb the vessel by heating it or pumping additional gas into it. As long as a unity's organization is maintained in the presence of such changes, they are termed *nondestructive*. In contrast, a structural change is *destructive* if it

destroys a unity's organization. For example, a pressure vessel may lose its integrity as a result of a corrosive chemical reaction with the gas it contains or because it has been heated to the point of melting. In this case, the unity disintegrates, losing its identity as an instance of the abstract class of which it was previously a member.

So far, so traditional. The crucial deviation of MV's framework from a traditional perspective on science is its recognition of the role of the observer within the framework. On MV's view, a unity such as a pressure vessel is not given a priori. Rather, it arises as a distinction made in language among a community of observers. Although by no means unique in modern science (consider quantum mechanics), making explicit the role of the observer in science takes out an enormous explanatory loan. Fortunately, MV proceed to fully repay this loan by systematically constructing an account of the observer within their framework. This construction proceeds in two steps. First, MV provide a theory of living systems. Then, they provide a theory of cognition built upon this biological foundation.

Given MV's perspective on science, it should not be surprising that their account of life consists not of a list of necessary properties or components, but rather a specification of the organization that is unique to living systems. MV call this organization *autopoietic*, which literally means "self-producing." The basic idea is that an autopoietic system consists of a network of processes that produce components whose interactions serve to generate and maintain the very network of processes that produced them. The canonical example of an autopoietic system is a living cell, which consists of a set of spatiotemporally organized molecular components that participate in biochemical reactions that produce those very components and maintain their spatiotemporal organization. Loosely speaking, a living cell is like a pressure vessel in which the reactions between the gas and the container are such as to preserve the vessel's integrity even in the face of natural degradation and external assault. Of course, a living system can actively maintain its autopoietic organization in the face of structural change only within limits. If those limits are exceeded, autopoiesis is lost and the unity disintegrates.

MV call the set of all interactions that an autopoietic system can participate in without loss of identity its *cognitive domain*. This is a somewhat unusual and controversial use of the term "cognitive" (a more neutral term would be "domain of interactions"). Nevertheless, MV's use of this term is grounded in the idea that the only way an entity can "know" a given environmental configuration is for that entity to undergo a nondestructive change in structure as a result of interacting with that configuration without disintegration. In this way, perception is conceptualized as a perturbative rather than an instructive interaction, because how an entity's structure changes in response to a given environmental perturbation depends as much on the entity's structure at the time of the interaction as it does on the structure of the perturbation itself. Thus, to MV, cognition is the process of identity-preserving action within an environment. An interesting consequence of this definition is that, strictly speaking, all living systems are cognitive systems; nervous systems are not essential for cognition. However, by introducing the potential for massive structural variation in a metabolically nonessential way, nervous systems obviously greatly expand the range of perturbations that an entity can compensate for without loss of identity.

As long as an entity persists in some environment, its structural changes must necessarily exhibit a congruence with those of that environment; if they did not, the entity would disintegrate. Likewise, the structural changes of the environment depend on both its own internal dynamics and the perturbations induced by the entities with which it interacts. MV call this process of congruent structural change *structural coupling*. When an entity's environment contains other entities, then they can become structurally coupled with one another, with the actions of each serving as a source of perturbations to the others. As long as they participate in this interlocking chain of mutual perturbations, these entities generate a new *consensual domain* of *communicative interactions* in which each orients the others to future possibilities for action within their respective cognitive domains.

In MV's framework, communicative interactions within a consensual domain form the basis for linguistic interactions that ultimately make possible a community of observers who can distinguish in

language the simple and composite unities that science tries to explain [14]. However, MV's account of language is beyond the scope of this article. What we will attempt to do here is to concretize the notions of cognitive domain and structural coupling by analyzing in depth their application to a simple model.

3   Glider Autopoiesis

The game of Life (GoL) is a two-dimensional cellular automaton invented by John Conway and popularized by Martin Gardner in the pages of *Scientific American* [2, 11, 21]. If we define s_{ij} to be the binary state of the (ith, jth) cell in an unbounded rectangular lattice $\mathcal{L}$ and let Σ_{ij} give the number of 1s in that cell's Moore neighborhood (the eight cells directly surrounding it), then the individual cell update rule for GoL can be written as

$$s_{ij}^{(k+1)} = \phi\left(s_{ij}^{(k)}, \Sigma_{ij}^{(k)}\right) \equiv \begin{cases} 1 & \text{if } \left(s_{ij}^{(k)} = 1 \wedge \Sigma_{ij}^{(k)} = 2\right) \vee \Sigma_{ij}^{(k)} = 3 \\ 0 & \text{otherwise} \end{cases}$$

where k is a nonnegative integer and $\wedge$ and $\vee$ denote conjunction and disjunction, respectively. Further, let $f : \mathcal{L} \rightarrow \mathcal{L}$ be a function that updates the entire lattice by simultaneously applying ϕ to every cell in parallel. We will use the notation f^k to denote the kth iteration of this function, and the notation $\mathcal{L}_{i,j}^{(k)}$ with bounds on i and j to denote the state of a subset of the lattice at iteration k. Thus, GoL is a discrete-time, discrete-space dynamical system.

When we as external observers watch a randomly initialized GoL lattice evolve over time, we will typically witness a great deal of complicated activity, which quickly settles down into a much sparser set of localized spatiotemporal patterns that persist, propagate, and interact. The term *glider* in GoL is used to refer to particular patterns such as those shown at the top of Figure 1. A glider is the simplest instance of a general class of patterns called *spaceships*, which are periodic structures that move through the lattice at some velocity. As a spatiotemporal pattern, a glider can be distinguished by the following three properties: (1) It is a nonempty 3×3 pattern (2) that repeats every four updates
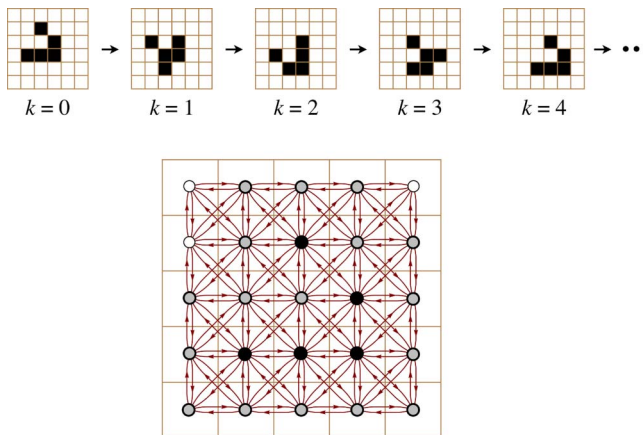

Figure 1. A glider in the game of Life. Top: Four iterations in the movement of a glider. Bottom: The lattice cell states constituting the $k = 0$ glider configuration (black nodes) depend upon the surrounding lattice cell states in their Moore neighborhoods (gray nodes) in order to make the proper transition to the $k = 1$ configuration.

(3) shifted by 1 cell in both the horizontal and vertical directions. Thus, we can say that a glider centered at location (i, j) exists in $\mathcal{L}$ at iteration k if the following set of equations is satisfied:

$$\mathcal{L}_{i+\Delta i, j+\Delta j}^{(k)} = f^4 \left(\mathcal{L}_{i+\Delta i \pm 1, \; j+Dj \pm 1}^{(k)} \right), \qquad -1 \leq \Delta i, \Delta j \leq 1$$

The four different sign assignments correspond to movement in the four possible diagonal directions ↗, ↘, ↙, and ↖. For each of the possible sign assignments, the resulting nine equations can be completely written out in terms of cell states at iteration k and nested applications of ϕ. There are nine cells in a 3 × 3 block and thus only $2^9 = 512$ possible 3 × 3 patterns. By checking all of them, we find that there are a total of 16 nonempty patterns that satisfy these equations, all of which are rotations or reflections of the $k = 0$ and $k = 1$ patterns shown at the top of Figure 1. It is important to emphasize that this spatiotemporal definition of a glider only makes sense against a local background of inactive cells. If any of the patterns in Figure 1 were surrounded by a field of active cells, not only would the resulting configuration be unrecognizable as a glider by an external observer, but it would fail to propagate or otherwise behave as one.

In a previous article [1], I suggested that spatiotemporal patterns such as gliders can also be interpreted as self-producing systems within a GoL lattice and thus may provide a simple model of a physical autopoietic system. The ideal way to test this proposal would be to provide a formal characterization of the organization of a glider and then examine the extent to which this organization satisfies MV's definition of autopoiesis. This is the focus of ongoing work. For the purposes of this article, I will simply describe the autopoietic formulation of a glider sketched in [1] so that we can focus here on exploring the cognitive implications of this formulation.

From an autopoietic perspective, we can think of a glider as consisting of configurations of components (the states of the underlying lattice cells) that participate in a network of processes (the GoL update rules acting over four iterations through the overlapping Moore neighborhoods of these components) that regenerate the very configurations of components whose interactions are necessary to maintain that network. Although the term "glider" is normally only applied to particular configurations of ON cells, it is easy to see that this definition is incomplete from an autopoietic perspective. In order for the ON cells (black nodes in the bottom of Figure 1) to undergo the transitions necessary for the preservation of a glider, all of the other cells in their Moore neighborhoods (gray nodes) must be OFF. Because the states and transitions of these OFF cells codetermine and are codetermined by the ON cells, they must be considered a fundamental part of a glider's constitutive network of processes. In addition, they form a kind of spatial boundary that delineates a glider from its surroundings. Thus, in this article, a glider will be defined as the 16 possible rotations and reflections of the two canonical patterns shown on the left side of Figure 2. We will use $\mathcal{G}$ to refer to this set of possible glider configurations.

For later reference, we introduce the naming convention Φ_θ^χ for these glider configurations. There are two basic forms Φ: W (the wedge shape) and R (the rocket shape). Each form can have four different orientations θ: 0, $\pi/2$, π, and $-\pi/2$. In addition, each orientation can occur with two different chiralities χ: R (right-handed) and L (left-handed). One interesting observation to make is that these three features—form, orientation, and chirality—are emergent properties of a glider; they do not exist at the level of the individual cells that constitute it. Higher-level glider features, such as the direction of glider movement, can be derived from these fundamental properties.

4 Cognitive Domain

In an empty environment, gliders undergo an endlessly repeating sequence of state changes determined by their own intrinsic dynamics (Figure 1). However, a nonempty environment can perturb
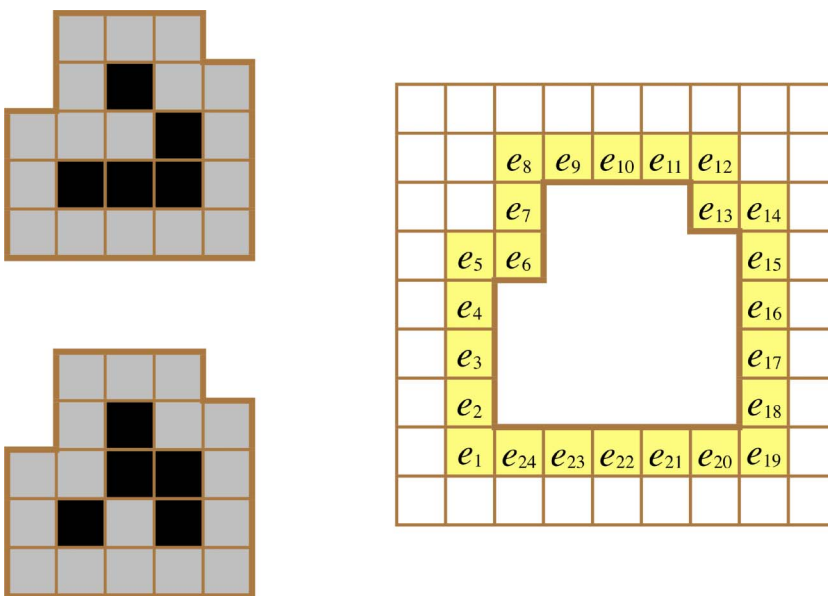
Figure 2. The canonical W (left top) and R (left bottom) glider states and a labeling of the cells in their immediate environment (in the electronic version, the yellow lattice cells are at the right).

this sequence in various ways. Recall that in MV's framework, an entity's cognitive domain is the set of all interactions in which it can participate without loss of identity. The goal of this section is to completely characterize the cognitive domain of a glider.

There are 24 environment cells immediately surrounding a glider (Figure 2, right). We will call this set of 24 cells a glider's 1-environment and postpone a more general discussion of larger environments until Section 7. Since each of the cells in the 1-environment can be either 0 or 1, there are 2^{24} possible perturbations $\mathcal{P}$ that a glider can undergo. For reasons that will become clear shortly, we will refer to these as microperturbations. The consequences of each microperturbation depend on the glider form, but not on its orientation or chirality (since the effect will be the same after the appropriate rotation or reflection of both the glider and its environment). Thus, we can impose all 2^{24} possible microperturbations on each of the two canonical W_0^R and R_0^R configurations in an otherwise empty environment and use symmetry to fill in the consequences for the remaining 14 configurations.

The results of this exhaustive microperturbation study are shown as an interaction graph in Figure 3. Of all perturbations, 99.57% are lethal for the W form and 99.87% are lethal for the R form. The remaining perturbations induce transitions from one glider state to another and thus preserve the glider's identity while changing its structure. These perturbations fall into classes $\mathcal{C}$ that we will refer to as macroperturbations and that we will label by the color of the corresponding arc in Figure 3. Thus, $\mathcal{C} = \{\text{BLACK, BROWN, BLUE, ORANGE, GRAY, GREEN}\}$. Two microperturbations fall into the same macroperturbation class if they perturb the same initial glider state into the same subsequent glider state. For example, all microperturbations that transform a W_0^R configuration into a $W_{\pi/2}^L$ configuration fall into the BLUE class. The BLUE class also contains all microperturbations that produce an equivalent transformation under rotational and/or reflectional symmetry. For example, all microperturbations that produce a transformation from a $W_{\pi/2}^R$ configuration to a W_π^L configuration (which are just $\pi/2$ rotations of the W_0^R and $W_{\pi/2}^L$ configurations, respectively) are also included in the BLUE class.

Note that the BLACK and GRAY macroperturbations are special because they include the 0-perturbation, that is, they are equivalent to the transition that a glider state would undergo in the absence of any environmental perturbation. We call perturbations in these two classes *null perturbations* because

they have no effect on the glider. If we focus on only the BLACK and GRAY arcs in Figure 3, we observe four distinct BLACK → GRAY → BLACK → GRAY 4-cycles, corresponding to unperturbed gliders moving along the four diagonal directions marked in red, in the electronic verison. Perturbations in the other classes trigger transitions within or between these different 4-cycles.

We can think of this interaction graph as defining a new parameterized function $F : \mathcal{G} \times \mathcal{C} \rightarrow \mathcal{G}$ that maps one glider state to another depending upon the class of perturbation applied. For example, we can say that $F(\mathrm{W}_0^{\mathrm{R}};$ ORANGE$) = \mathrm{R}_{\pi/2}^{\mathrm{R}}$. Because encountering some classes of perturbations in some states is lethal for a glider, F is only a partial function. In order to make it a total function, we define an additional state $\varnothing$ to represent the absence of a glider, extend the domain of F to $\mathcal{G} \cup \{\varnothing\} \times \mathcal{C} \rightarrow \mathcal{G} \cup \{\varnothing\}$, and define $\varnothing$ to be absorbing, that is, $F(\varnothing; c) = \varnothing$ for all possible perturbation classes $c \in \mathcal{C}$ (once a glider is dead, its stays dead).

Note that, although the function F on glider states is ultimately grounded in the lattice update function f on cell states defined in Section 2, the two are not identical. The function f describes the microdynamics of the constituents of a glider, and the cognitive map F describes the macrodynamics of glider entities as they undergo transformations in response to environmental perturbations. Whereas f applies to any lattice configuration, the function F can only be applied to a glider when one is actually present in the lattice. Thus, the cognitive domain can be interpreted as a new emergent domain of description.
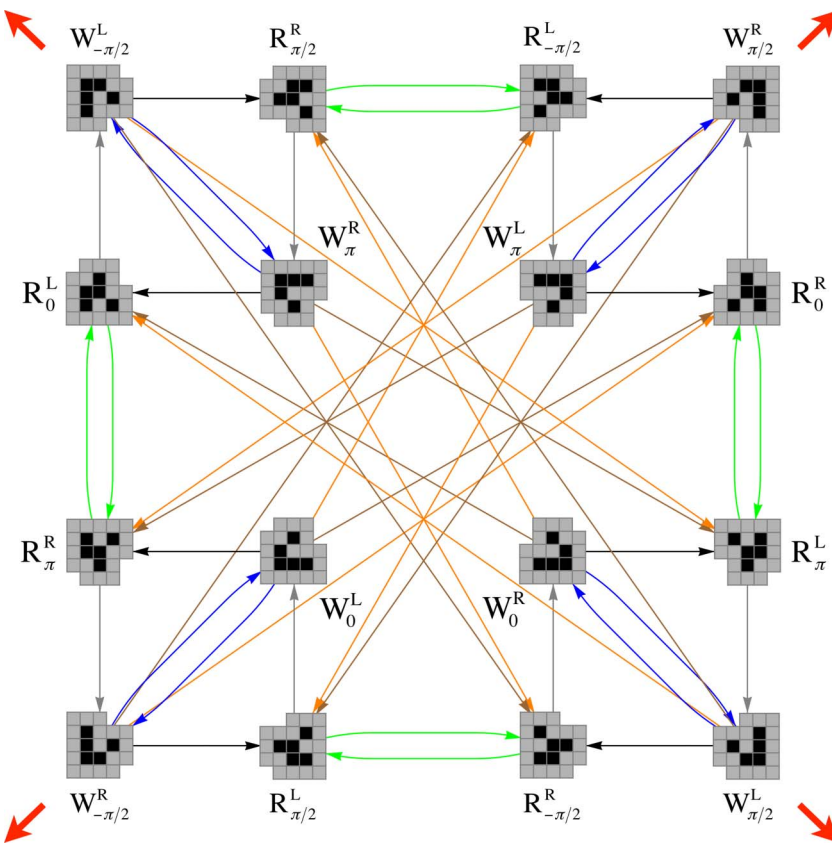
Figure 3. The cognitive domain of a glider. Each glider configuration is labeled as described in the main text, and transitions between them are colored according to the class of the perturbations that produce that transition. In the electronic version, red arrows indicate the direction of motion of each of the four BLACK → GRAY → BLACK → GRAY 4-cycles through the lattice.

Because the glider organization is preserved by all perturbations within its cognitive domain, we can also characterize the various classes of perturbations by the glider properties that they maintain invariant:

1. Form is preserved by BLUE and GREEN perturbations only; all other perturbations swap it.

2. Chirality is preserved by GRAY and ORANGE perturbations only; all other perturbations swap it.

3. Orientation is preserved by BROWN perturbations only; GRAY, BLUE, and ORANGE perturbations rotate by $\pm\pi/2$, whereas BLACK and GREEN perturbations rotate by π.

4. BLACK perturbations have no invariants, simultaneously swapping form and chirality and rotating by π.

Finally, the direction of glider motion is preserved by BLACK, GRAY, and BLUE perturbations, whereas BROWN and ORANGE perturbations reverse the direction of motion, and GREEN perturbations can reflect about either the horizontal or the vertical axis, depending on the chirality and orientation of the R form to which they are applied.

5 Glider Epistemology

Each macroperturbation class contains many microperturbations. What features do all microperturbations in a given macroperturbation class share? How does a glider distinguish between microperturbations falling into different macroperturbation classes, but fail to distinguish between microperturbations in the same class? Recall that, according to MV, for an entity to "know" a given environmental configuration is for it to undergo a nondestructive change in structure as a result of interacting with that configuration without disintegration. Because we as scientific observers stand outside the GoL lattice, we can simultaneously consider the relationship between all possible environmental configurations and all possible glider states. Therefore, we can characterize the sets of distinctions that a glider can and cannot make about the state of its immediate environment and how those distinctions are grounded in the particular way that a glider maintains its identity. In this sense, we are in a position to completely understand the "epistemology" of a glider.

By simply enumerating the effect of each microperturbation from the exhaustive perturbation study described in the previous section, we find that, of the $2^{24} = 16,777,216$ possible microperturbations to a glider, 16,683,428 destroy it. Thus, the vast majority of immediate environmental configurations are fundamentally unknowable by a glider. The 93,788 remaining configurations are knowable depending on the glider's form, with the following breakdown. For the W form, there are 8,320 BLACK perturbations, 8,192 BROWN perturbations, 21,112 BLUE perturbations, and 34,816 ORANGE perturbations, for a total of 72,440 survivable perturbations. For the R form, there are 15,456 GRAY perturbations and 7,040 GREEN perturbations, for a total of 22,496 survivable perturbations. Interestingly, 1,148 perturbations are shared by both the BLACK and GRAY classes; a glider's response to these perturbations depends entirely on the state it is in when they are encountered. Otherwise, all of these classes are disjoint. For example, any BLUE perturbation is lethal to the R form, and any GREEN perturbation is lethal to the W form.

In von Uexküll's terms [30], these classes form a glider's *Umwelt*. From the 2^{24} environmental configurations that we can distinguish, a glider can distinguish only the six classes in $\mathcal{C}$: Microperturbations from within the same macroperturbation class produce the same glider state and thus the same response to a subsequent perturbation. Furthermore, the structure of these classes may appear nonsensical to us. We normally describe biological organisms interacting with their environments in terms of such categories as "terrain," "food," "predator," and so on, but no such categories are apparent here. However, our task is not to evaluate the significance of these classes

to us, but rather to understand how a glider's organization and structure make them significant to it. We do this by characterizing the structure of each class and showing how this structure derives from the constraints that must be satisfied for a glider to maintain its existence in response to a perturbation.

The structure of each class is summarized in Figure 4. Each image in this figure shows the initial glider state and a representation of the constraints required to make the named transition occur. The following environmental cell coloring convention is used. Black and white environmental cells are constrained to be either ON or OFF, respectively, in order for the given transition to take place. Green environment cells are partially constrained; although no single pattern of activity is required for these cells, only a subset of the possible patterns is permissible. Yellow environmental cells are completely unconstrained. Interestingly, we can see that different transitions are sensitive to different combinations of environmental cells; in effect, different *sensory surfaces* emerge at different locations around the glider, depending on its state and the immediate situation. Furthermore, it turns out that nonlocal integration across these sensory surfaces is required for all of the nonnull transitions: At least two nonadjacent patterns of ON cells are necessary for these transitions to take place.

For a BLACK transition from a W-form glider to take place, environment cells e_6, e_{13}, and $e_{20}-e_{24}$ must be OFF. Cells e_1 and e_8-e_{12} can take on any value. Finally, cells e_2-e_7 must take on one of 13 patterns, and cells $e_{14}-e_{19}$ must take on one of 10 patterns. This gives a total class membership of $2^6 \times 13 \times 10 = 8{,}320$ perturbations.

For a BLUE transition from a W-form glider to take place, cell e_6 must be OFF and cell e_{13} must be ON. Cells e_1 and e_8 can take on any value. Finally, cells e_2-e_{13} must take on one of 13 patterns, and the remaining 15 green cells must take on one of 406 patterns. This gives a total class membership of $2^2 \times 13 \times 406 = 21{,}112$ perturbations.

For a BROWN transition from a W-form glider to take place, cell e_6 must be ON and cells e_{13} and $e_{20}-e_{24}$ must be OFF. Cells e_1-e_5 and e_7-e_{12} can take on any value. Finally, cells $e_{14}-e_{19}$ must take on one of four possible patterns. This gives a total class membership of $2^{11} \times 4 = 8{,}192$ perturbations.

For an ORANGE transition from a W-form glider to take place, cell e_6 must be ON and cell e_{13} must be OFF. Cells e_1-e_5 and e_7-e_{12} can take on any value. Cells $e_{14}-e_{24}$ must take on one of 17 different patterns. This gives a total class membership of $2^{11} \times 17 = 34{,}816$ perturbations.

For a GRAY transition from an R-form glider to take place, cells e_6, e_7, and e_{13} must be OFF. Cells e_1 and e_{24} can take on any value. Cells e_2-e_5 must take on one of 7 possible values and cells $e_{14}-e_{23}$ must take on one of 552 possible values. This gives a total class membership of $2^2 \times 7 \times 552 = 15{,}456$ perturbations.

Finally, for a GREEN transition from an R-form glider to take place, cells e_6, e_7, and e_{13} must be ON. Cells e_1-e_5 and e_{24} can take on any value. Cells e_8-e_{12} and $e_{14}-e_{23}$ must take on one of 110 patterns. This gives a total class membership of $2^6 \times 110 = 7{,}040$ perturbations.

Next, we illustrate the process of deriving the structure of a perturbation class from the constraints that must be satisfied for a glider to make the required state transition in response to perturbations from that class. Consider the BROWN transition shown at the top of Figure 5. To begin, we draw the initial glider state, with its ON cells in black, its OFF cells in gray, and its environment cells in yellow to indicate that they can in principle take on any value (bottom of Figure 5). We then superimpose a transparent image of the target glider state, with its ON cells in blue and its OFF cells in pink, in the proper position. Finally, for every cell in the target glider state, we fill in the number of ON cells in its Moore neighborhood. Some of these sums are known exactly, and others are known only symbolically because they depend on the unknown states of the environment. We wish to solve for the set of environmental states that allow this transition to take place. Note that, since they do not enter into the neighbor sums of any cell in the target glider states, we can already see that environment cells e_1-e_5 and e_7-e_{12} can take on any value without affecting a BROWN transition, which is why they remain yellow in the electronic version of Figure 4.
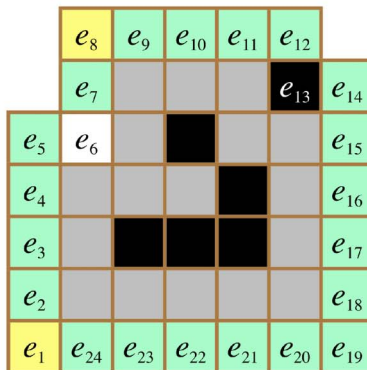
We next write down the full set of equations that must be satisfied for each cell in the initial state to have the correct value in the target state. Since we can immediately verify that each cell with a
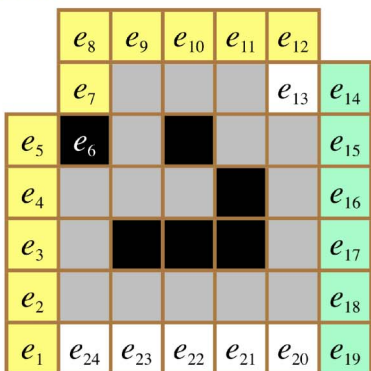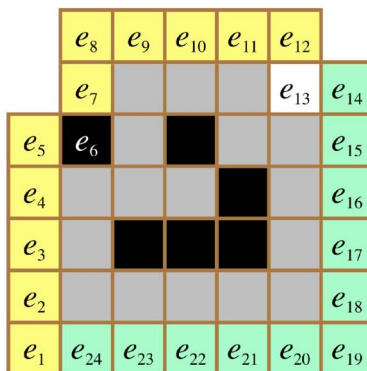
BLACK
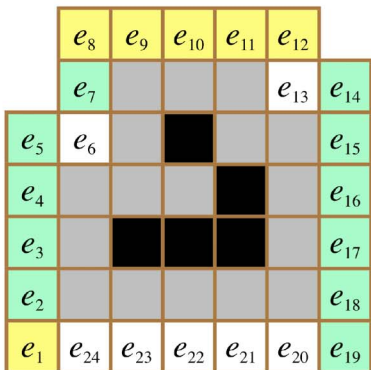
$$2^6 \times 13 \times 10 = 8{,}320$$
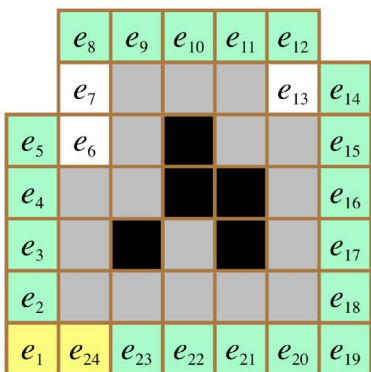
BLUE

$$2^2 \times 13 \times 406 = 21{,}112$$

BROWN

$$2^{11} \times 4 = 8{,}192$$
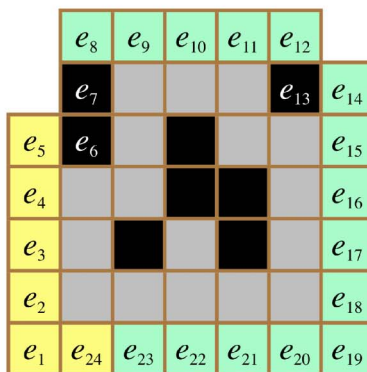
ORANGE

$$2^{11} \times 17 = 34{,}816$$

GRAY

$$2^2 \times 7 \times 552 = 15{,}456$$

GREEN

$$2^6 \times 110 = 7{,}040$$

Figure 4. The structure of the six perturbation classes distinguished by a glider. In the electronic version, in order for each of the named transitions to occur, environment cells shown in black and white must be ON and OFF, respectively, while green cells are partially constrained and yellow cells are unconstrained. The cardinality of each class is given below the corresponding diagram, with powers of 2 accounting for the unconstrained yellow cells and each remaining factor accounting for the allowable patterns of a distinct subset of partially constrained green cells.

numerical neighbor count will make the proper transition, we will not consider these states further. For the nonnumerical cell counts, we reason as follows. Consider the cell with neighborhood count $e_6 + 3$. This cell is OFF in the initial state and should remain OFF in the target state. The only way that this cell can fail to remain OFF is for this sum to equal 3. This gives rise to the inequality $e_6 + 3 \neq 3$. Similar reasoning (with conditionals arising for environment states that
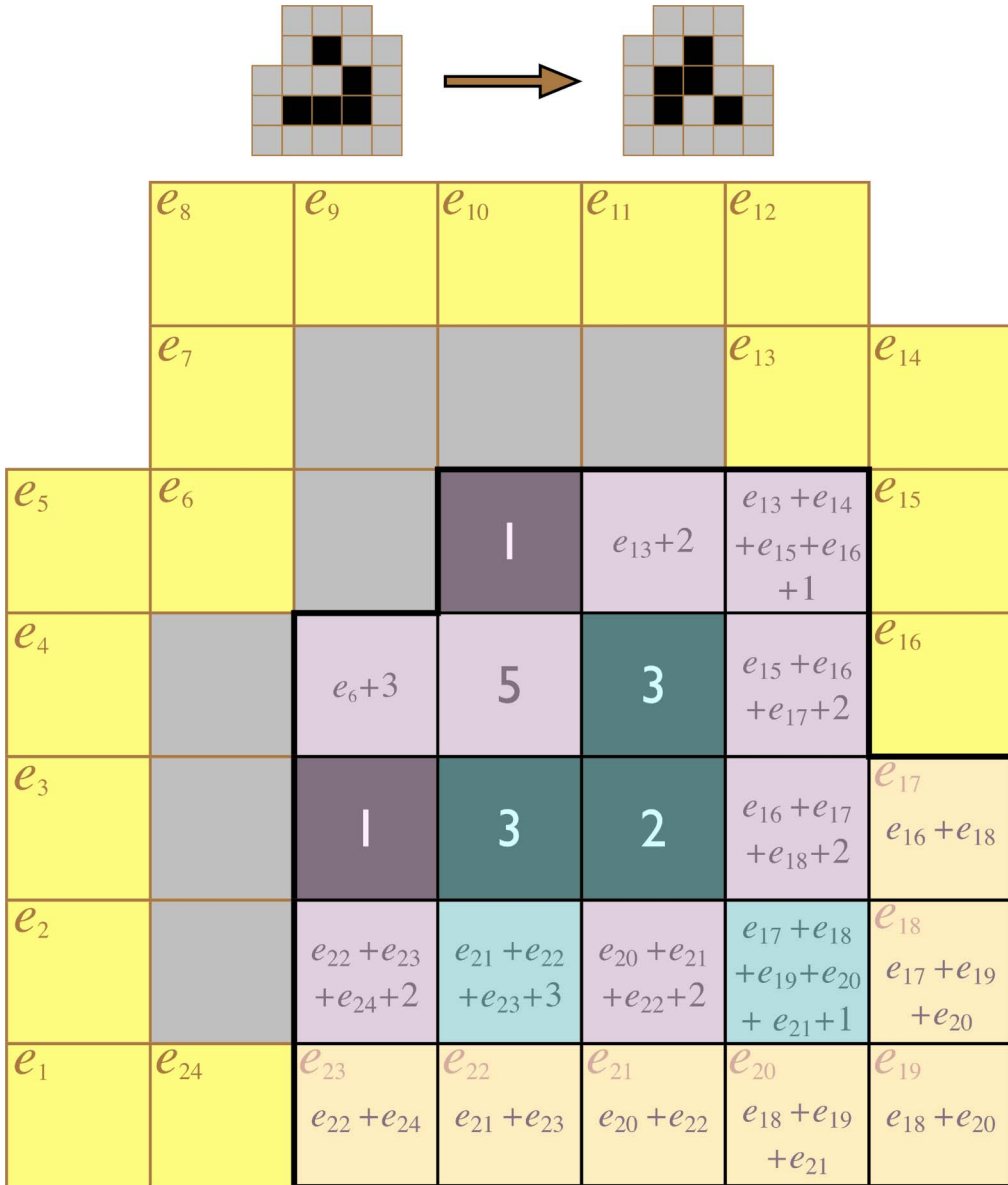
Figure 5. In the electronic version, neighborhood counts involved in determining the membership of the BROWN perturbation class. The original W state is shown in black and gray, surrounded by its immediate environment in yellow. The R state resulting from a BROWN perturbation is shown in its proper position as a transparent overlay, with its ON cells colored blue and its OFF cells colored pink. The neighborhood counts of each cell in the new R state are shown as an integer when they are known, and as a symbolic expression involving unknown environment states otherwise.

become part of the target glider because their initial states are unknown) gives rise to the following set of constraint equations:

$$e_6 + 3 \neq 3$$
$$e_{13} + 2 \neq 3$$
$$e_{13} + e_{14} + e_{15} + e_{16} + 1 \neq 3$$
$$e_{15} + e_{16} + e_{17} + 2 \neq 3$$
$$e_{16} + e_{17} + e_{18} + 2 \neq 3$$
$$e_{17} + e_{18} + e_{19} + e_{20} + e_{21} + 1 = 3$$
$$e_{20} + e_{21} + e_{22} + 2 \neq 3$$
$$e_{21} + e_{22} + e_{23} + 3 = 3$$
$$e_{22} + e_{23} + e_{24} + 2 \neq 3$$
$$e_{16} + e_{18} \neq \begin{cases} 3, & e_{17} = 0 \\ 2, 3, & e_{17} = 1 \end{cases}$$
$$e_{17} + e_{19} + e_{20} \neq \begin{cases} 3, & e_{18} = 0 \\ 2, 3, & e_{18} = 1 \end{cases}$$
$$e_{18} + e_{20} \neq \begin{cases} 3, & e_{19} = 0 \\ 2, 3, & e_{19} = 1 \end{cases}$$
$$e_{18} + e_{19} + e_{21} \neq \begin{cases} 3, & e_{20} = 0 \\ 2, 3, & e_{20} = 1 \end{cases}$$
$$e_{20} + e_{22} \neq \begin{cases} 3, & e_{21} = 0 \\ 2, 3, & e_{21} = 1 \end{cases}$$
$$e_{21} + e_{23} \neq \begin{cases} 3, & e_{22} = 0 \\ 2, 3, & e_{22} = 1 \end{cases}$$
$$e_{22} + e_{24} \neq \begin{cases} 3, & e_{23} = 0 \\ 2, 3, & e_{23} = 1 \end{cases}$$

Our next step is to simplify this set of equations as much as possible. Two trivial simplifications are immediately obvious. Since all environment cells must take on the value 0 or 1, the only way that the constraint $e_6 + 3 \neq 3$ can be satisfied is if $e_6 = 1$, which is why this cell is black in Figure 4. Likewise, the only way that the constraint $e_{13} + 2 \neq 3$ can be satisfied is if $e_{13} = 0$, and knowing this value removes e_{13} from the other equation in which it appears. Another simplification comes from the constraint $e_{21} + e_{22} + e_{23} + 3 = 3$, which can only be true if $e_{21} = e_{22} = e_{23} = 0$. This conclusion then triggers a chain of other simplifications that result in the conclusion that $e_{20} = e_{24} = 0$, which is why all of these cells are white in Figure 4. Finally, we can simplify the conditional inequalities by observing that the sum of the values of two environment cells can never equal 3, so that any assertion of this fact is redundant. Applying all of these simplifications leaves us with the following set of unconditional constraints for the values of e_{14}–e_{19}, which is why they are green in the electronic version of Figure 4:

$$e_{14} + e_{15} + e_{16} \neq 2$$
$$e_{15} + e_{16} + e_{17} \neq 1$$
$$e_{16} + e_{17} + e_{18} \neq 1$$
$$e_{16} + e_{17} + e_{18} \neq 3$$
$$e_{17} + e_{18} + e_{19} = 2$$

Solving this set of equations gives the four possible patterns that these cells can take on in a BROWN perturbation: $\{e_{14}, e_{15}, e_{16}, e_{17}, e_{18}, e_{19}\} = \{0, 1, 0, 1, 1, 0\}$, $\{0, 0, 1, 1, 0, 1\}$, $\{1, 1, 1, 1, 0, 1\}$, $\{1, 1, 1, 0, 1, 1\}$.
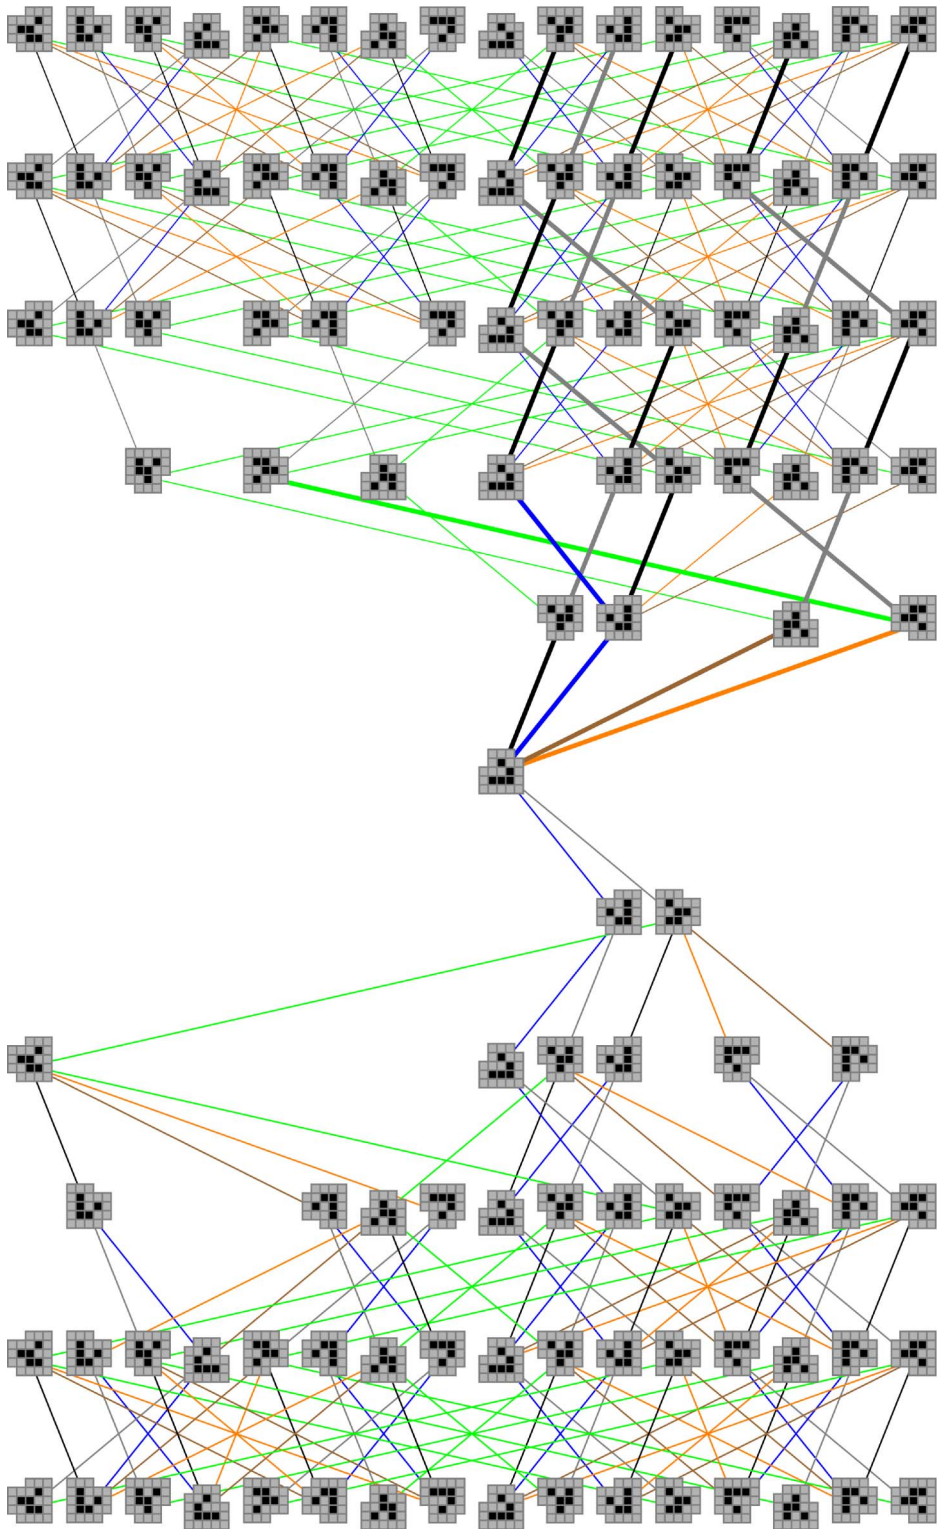
Figure 6. All behavioral trajectories through the canonical W state, from five steps in the past through five steps in the future. Thick lines indicate the subset of behavioral trajectories that result from structural coupling.

Similar calculations can be made for the other five perturbation classes. In each case, following the procedure just outlined, a set of explicit constraint equations can be derived that describe which subsets of environmental cells must take on fixed values, which can take on any of a constrained set of values, and which can take on any value in order for a given transition to take place. Furthermore, the process of deriving these constraint equations illuminates how they are grounded in the requirement that a glider maintain its identity in response to each class of perturbation.

6 Behavioral Trajectories

To this point, we have only considered the one-step effect of individual 1-environment perturbations on a glider. However, over time, a persisting glider will undergo a sequence of state changes as a result of the interaction between its own internal dynamics and the sequence of perturbations that it experiences. We will call such sequences *behavioral trajectories*, because an observer would interpret them as actions that the glider is taking (turning to the left, moving one step to the right, etc.). What can we say about the structure of such behavioral trajectories? In this section, we will characterize the dynamics of a glider driven by imposed sequences of 1-environment macro-perturbations. This will prepare us for an examination in Section 7 of how MV's concept of structural coupling can be applied to the situation where a glider and its environment are in a relationship of mutual perturbation.

Formally, this involves studying the dynamics of iterates of the cognitive map F, that is, $F^k(g_0; c_0,...,c_{k-1})$, where g_0 denotes an initial glider state drawn from $\mathcal{G}$, and $c_0,...,c_{k-1}$ denotes a sequence of perturbation classes, each of which is drawn from $\mathcal{C}$. Thus, for example, $F^2(W_0^R; \text{BLUE}, \text{ORANGE}) = R_0^L$, but $F^k(W_0^R; \text{BLUE}, \text{GREEN}, ...) = \varnothing$ for all $k > 1$. The latter fate can be avoided by applying only BLACK, BROWN, BLUE, or ORANGE perturbations to W-states and applying only GRAY or GREEN perturbations to R-states.

Since behavioral trajectories are just paths through the cognitive domain, it is useful to visualize the structure of these paths for a particular glider state by unrolling the interaction graph over time. Figure 6 shows all paths through the canonical wedge state W_0^R, from five iterations in the past to five iterations in the future. Each path through this layered graph represents a valid behavioral trajectory. An analogous graph can be constructed for any glider state in the cognitive domain. Note that paths can both diverge and converge, depending upon the sequence of perturbations encountered. Note also that iteration 4 includes all 16 glider states. Thus, there always exists a sequence of perturbations of length at most 4 that can take a glider from one state to any other state. An example of a transition that requires a minimum of four steps is $W_0^R \rightarrow W_\pi^L$, which can only be accomplished by a length-4 perturbation sequence such as {BLUE, BROWN, GREEN, GRAY}. An important consequence of this fact is that the path structure from iteration 4 to iteration 5 repeats indefinitely into the future (and similarly for the path structure from iteration −4 to iteration −5 into the past).

How a glider responds to a given perturbation depends on the state it is in at the time, which in turn depends at least in part on the sequence of perturbations that it has previously experienced. Thus, as a glider encounters and survives a particular sequence of perturbations, dependencies build up between the glider's state and that sequence. Our formalization of a glider's cognitive domain allows us to pose and answer a variety of interesting questions about the structure of these dependencies for a glider, both in forward and backward time.

The kinds of forward questions that we can address include the following:

(1) Given an initial glider state and a sequence of k perturbations, what state does it end up in? The answer to this question is simply $F^k(g_0; c_0,...,c_{k-1})$. For example, we noted above that $F^4(W_0^R; \text{BLUE}, \text{BROWN}, \text{GREEN}, \text{GRAY}) = W_\pi^L$.
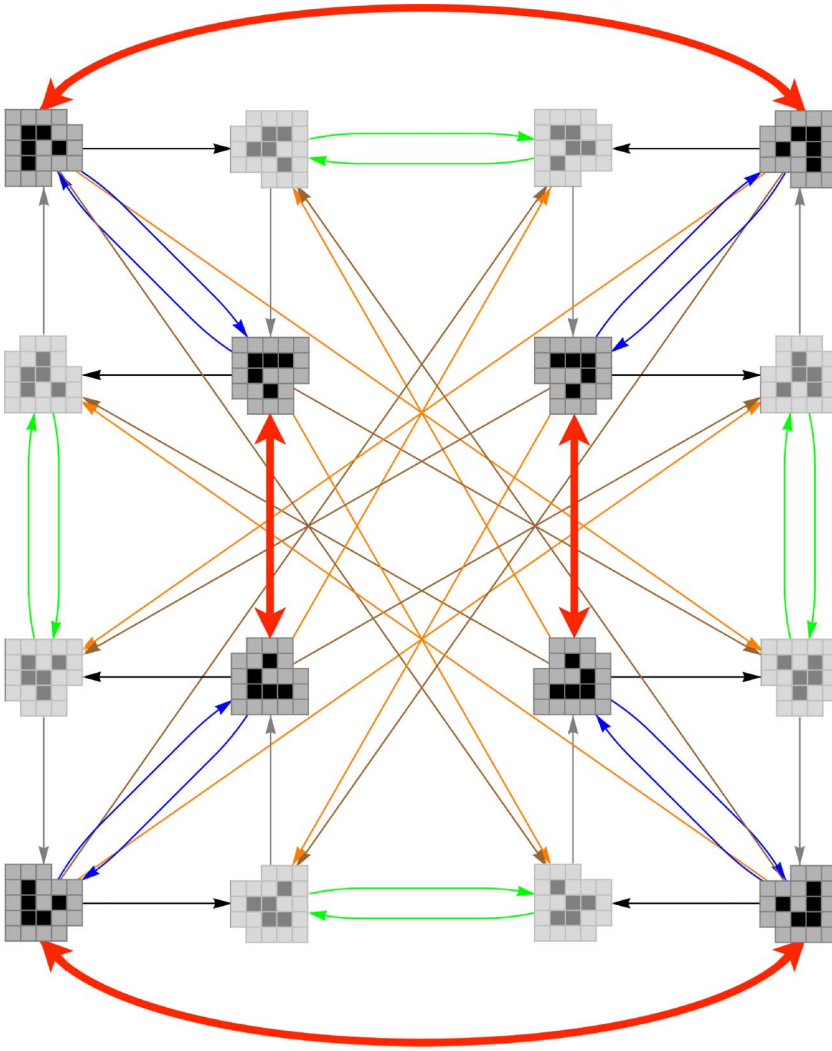
Figure 7. The effect of a {BLUE, BROWN, GREEN, GRAY} perturbation sequence on a glider's cognitive domain. In the electronic version, R-states are destroyed, and W-states are transformed as shown by the red arrows.

(2) Given a sequence of k perturbations, how does what state a glider ends up in depend on what state it started in? This question is answered by studying $F^k(g; c_0,...,c_{k-1})$ as a function of g. For example,

$$F^4(g; \text{BLUE, BROWN, GREEN, GRAY}) = \begin{cases} W^L_\pi & \text{if } g = W^R_0 \\ W^L_{-\pi/2} & \text{if } g = W^R_{\pi/2} \\ W^L_0 & \text{if } g = W^R_\pi \\ W^L_{\pi/2} & \text{if } g = W^R_{-\pi/2} \\ W^R_\pi & \text{if } g = W^L_0 \\ W^R_{-\pi/2} & \text{if } g = W^L_{\pi/2} \\ W^R_0 & \text{if } g = W^L_\pi \\ W^R_{\pi/2} & \text{if } g = W^L_{-\pi/2} \\ \varnothing & \text{otherwise} \end{cases}$$

This sequence reflects W-form gliders about either their horizontal or vertical axes, depending on their orientation (red arrows in the electronic version of Figure 7) while destroying all R-form gliders.

(3) Given an initial glider state, how does what state a glider ends up in depend on what sequence of perturbations it encounters? This question is answered by studying $F^k(g_0; C^k)$ as a function of the possible length-k perturbations C^k. For example,

$$F^2\big(W_0^R; C^2\big) = \begin{cases} W_{\pi/2}^L & \text{if } C^2 = \{\text{BLACK}, \text{GRAY}\} \\ R_0^R & \text{if } C^2 = \{\text{BLACK}, \text{GREEN}\} \\ W_0^R & \text{if } C^2 = \{\text{BLUE}, \text{BLUE}\} \\ R_{\pi/2}^R & \text{if } C^2 = \{\text{BLUE}, \text{BROWN}\} \\ R_0^L & \text{if } C^2 = \{\text{BLUE}, \text{ORANGE}\} \\ W_{-\pi/2}^L & \text{if } C^2 = \{\text{BROWN}, \text{GRAY}\} \\ R_\pi^R & \text{if } C^2 = \{\text{BROWN}, \text{GREEN}\} \\ W_\pi^R & \text{if } C^2 = \{\text{ORANGE}, \text{GRAY}\} \\ R_{-\pi/2}^L & \text{if } C^2 = \{\text{ORANGE}, \text{GREEN}\} \\ \varnothing & \text{otherwise} \end{cases}$$

Of course, all of these forward questions are just special cases of the general question of how the state a glider ends up in at a given time depends on where it starts and the sequence of perturbations it encounters, that is, $F^k(g, C^k)$ for arbitrary k, g, and C^k, which in turn is really just a question about the structure of all possible forward paths through the cognitive domain.

The kinds of backward questions that we can address include the following:

(1) What initial state g leads to a given final state g_f after a given sequence of perturbations, that is, $\{g \,|\, F^k(g; c_0,...,c_{k-1}) = g_f\}$? For example, $\{g \,|\, F^4(g, \{\text{BLUE}, \text{BROWN}, \text{GREEN}, \text{GRAY}\}) = W_\pi^L\} = \{W_0^R\}$.

(2) What sequence of perturbations leads from a given initial state to a given final state, that is, $\{C^k \,|\, F^k(g_0; C^k) = g_f\}$? For example, $\{C^4 \,|\, F^4(W_0^R; C^4) = W_\pi^L\} = \{\{\text{BLUE}, \text{BROWN}, \text{GREEN}, \text{GRAY}\}, \{\text{BLACK}, \text{GREEN}, \text{GRAY}, \text{BLUE}\}\}$. Once again, these are just special cases of the general question of how knowing a glider's state in the present constrains its possible past states and perturbation sequences, that is, $\{g, C^k \,|\, F^k(g, C^k) = g_f\}$, which is really just a question about the structure of all possible backward paths through the cognitive domain.

Thus, given F, it is relatively straightforward to answer any particular instance of these forward and backward questions. But can we say anything more general about the structure of dependencies in a glider's behavioral trajectories? In fact, information theory [4] can be used to quantify the information that observing a given glider state now gives us about the glider's possible pasts and futures. In order to do so, we must be explicit about the universe of possible perturbations that we wish to study. There are four possibilities: (1) the set of all sequences of microperturbations; (2) the subset of all sequences of microperturbations that preserve glider identity; (3) the set of all sequences of macroperturbations; (4) the subset of all sequences of macroperturbations that preserve glider identity. We have mostly focused on (4) in this article, and we will continue to do so here by assuming a uniform distribution over nonlethal sequences of macroperturbations. But it is important to understand that other choices are possible and that different choices will in general produce different results.

As an example of the application of information theory to characterize the dependence between a glider's state changes and the sequence of environmental perturbations it encounters, suppose we

know that a glider began in state W_0^R and we want to quantify the information that a measurement of its state after k iterations carries about the sequence of perturbations it has encountered in the interim. This is given by the mutual information $I(C^k; F^k(W_0^R, C^k))$, where the mutual information between two random variables X and Y can be defined in terms of their entropies and conditional entropies as $I(X; Y) = H(Y) - H(Y|X)$. From this point on we will normalize everything by $H(X)$ (which just equals $\log |C^k|$ in our case, since we are assuming a uniform distribution over nonlethal paths of length k). Also, because Y is a function of X in our case, $H(Y|X) = 0$ and we are left with $I(C^k; F^k(W_0^R, C^k)) = H(F^k(W_0^R, C^k))$. If we plot this information as a function of k (blue curve in the electronic version of Figure 8), we see that the initially complete information about perturbation paths begins to fall steadily toward a limiting form (yellow curve) as paths begin to converge at $k = 3$. This limiting form is given by $I(C^k; F^k(g, C^k))$, the path information available in a state measurement at iteration k when the initial state is unknown, which can be derived in closed form. Similar calculations can be done to characterize the informational relationships for any of the forward or backward questions described in this section.

7   Structural Coupling

The analysis in the previous section treated each 1-environment perturbation as if it were completely independent of the preceding one. Although we as external observers can impose such perturbation sequences on a glider, this is clearly not the case within GoL itself, since environmental states must follow the same underlying laws of "physics" as glider states. In addition, just as the environment perturbs the entity, so does the entity in turn perturb the environment. Recall that MV use the term *structural coupling* to refer to the mutual selection that occurs between a surviving entity and its environment of paths of structural change from among the possible paths of each. How might we apply this concept of structural coupling to the interactions that gliders can undergo with their environments?

Let us begin by fully analyzing the simplest possible scenario. Suppose that we impose one of the allowable 1-environment perturbations on a particular glider state as an initial condition in an otherwise empty lattice, and then subsequently allow the state of the environment to evolve "naturally" rather than manually imposing additional allowable perturbations. This scenario is the natural extension of the one studied in the previous section. Examples of it are shown in Figure 9.
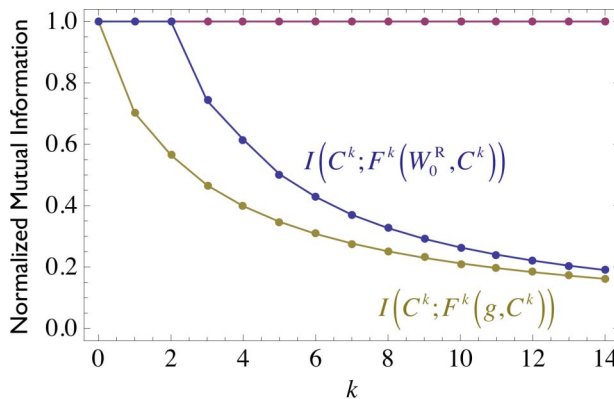

Figure 8. The information about the sequence of environmental perturbations a glider has encountered that is available in a measurement of that glider's state at iteration k. In the electronic version, the blue curve averages over all macroperturbation class sequences from the canonical W state, whereas the yellow curve averages over all macroperturbation class sequences from all possible initial states. Note that the blue curve asymptotically approaches the yellow one. The magenta curve considers only structurally coupled paths from the canonical W state (thick lines in Figure 6). Note that the state of a structurally coupled glider retains perfect information about the sequence of perturbations that it has encountered.
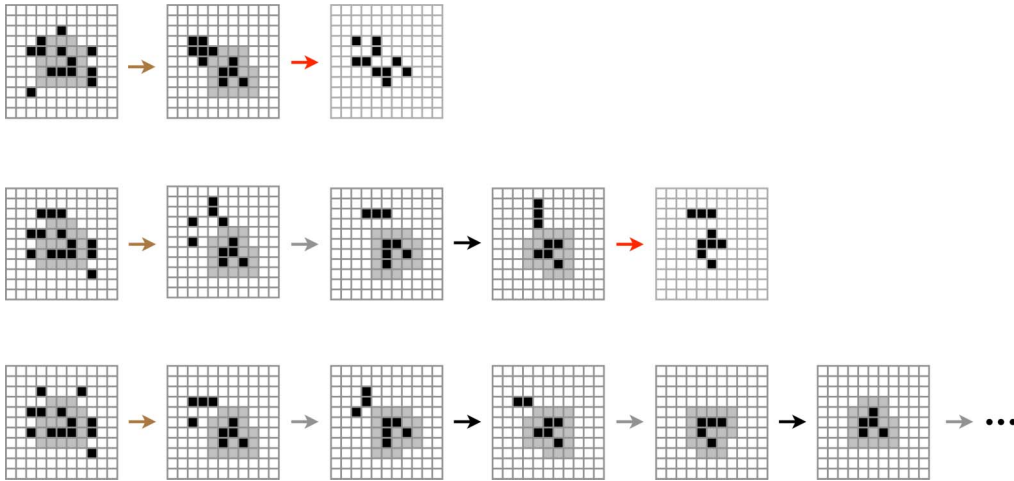
Figure 9. Three examples of structurally coupled interactions that, although they all begin with a BROWN perturbation, evolve in very different ways.

Note that in order to perform this analysis, we need to switch our focus from macroperturbation classes to a consideration of microperturbations, because two microperturbations in the same class will in general evolve to different subsequent environment states. Figure 9 shows three examples of initial BROWN perturbations with very different eventual fates. The glider in the top example survives only one step, in the middle example survives three steps, and in the third example survives indefinitely.

In order to provide a broader picture of the consequences of structural coupling in this scenario, let us revisit the macroperturbation pathway graph shown in Figure 6. Recall that this graph shows all the possible paths to depth 5 in the future and the past for the canonical wedge state. Of the 208 possible forward paths from W_0^R shown, only five survive to level 5 in the presence of structural coupling, with one surviving only to level 2 (thick lines in Figure 6). Note that, beyond level 2, only null macroperturbation paths survive (alternating gray and black lines). Note also that these structurally coupled paths do not converge. This implies that, unlike in the case of unconstrained perturbation sequences, there is no loss of information about the past for structurally coupled paths. This can be seen in Figure 8 (magenta line in the electronic version), where a measurement of the glider state at any time provides perfect information about the sequence of macroperturbations the glider has experienced to that point. This result is perhaps somewhat surprising and would certainly not be expected to be true for arbitrary environments. However, it is true for the kinds of interactions that develop between a glider and all possible 1-environments in an otherwise initially empty lattice.

As mentioned above, macroperturbations do not tell the entire story in the presence of structural coupling. Just because some initially BROWN perturbations survive until level 5, that does not mean they all do, as we have already seen in Figure 9. Indeed, of the 72,440 microperturbations of W_0^R that survive to level 1, only 4,458 survive to level 2, only 2,362 survive to level 5, and only 2,241 survive beyond level 10. A similar analysis can be performed on R_0^R, with similar results. Thus, structural coupling strongly constrains the interactions that a glider will engage in without loss of identity. It selects from all possible trajectories of glider–environment dynamics only those that are consistent with a glider's ongoing existence—or, from the perspective of the glider, it selects a subset of the possible paths through its cognitive domain.

To reiterate, the specific results we have described so far in this section apply only to the simplified scenario of structural coupling between a glider and all possible 1-environment initial states in an otherwise empty lattice. Obviously, larger sets of cells can impact a glider over longer periods of time, and we would ideally also like to examine structural coupling in such larger environments. Unfortunately, this is impractical to do in any exhaustive way. The 24 cells of the 1-environment are

surrounded by 32 cells of what we can call the 2-environment, and so on. In general, the *l*-environment contains $8l + 16$ cells, and therefore the total $(1,l)$-environment contains $\sum_{i-1}^{l} 8i + 16 = 4l^2 + 20l$ cells with 2^{4l^2+20l} possible initial environment states for a given glider state. Even in finite lattices with periodic boundary conditions, this number of initial states very quickly becomes computationally infeasible.

A preliminary study of a sample of the (1,2)-environments has revealed a variety of interesting behavioral trajectories, several examples of which are shown in Figure 10. The top example undergoes a {BLUE, BROWN, GRAY, BLUE, BLUE} sequence of perturbations before disintegrating. This sequence of perturbations drives the glider's state along a path through its cognitive domain (solid bold arcs in the graph at the bottom of Figure 10) that is very different from the path it would follow in an empty environment (dashed arcs in graph). The middle example undergoes a {BLUE, BLUE, ORANGE} sequence before entering into a {GRAY, BLACK, GRAY BLACK, …} null perturbation loop that persists forever, leaving behind a static square pattern called a block. The bottom example begins on a
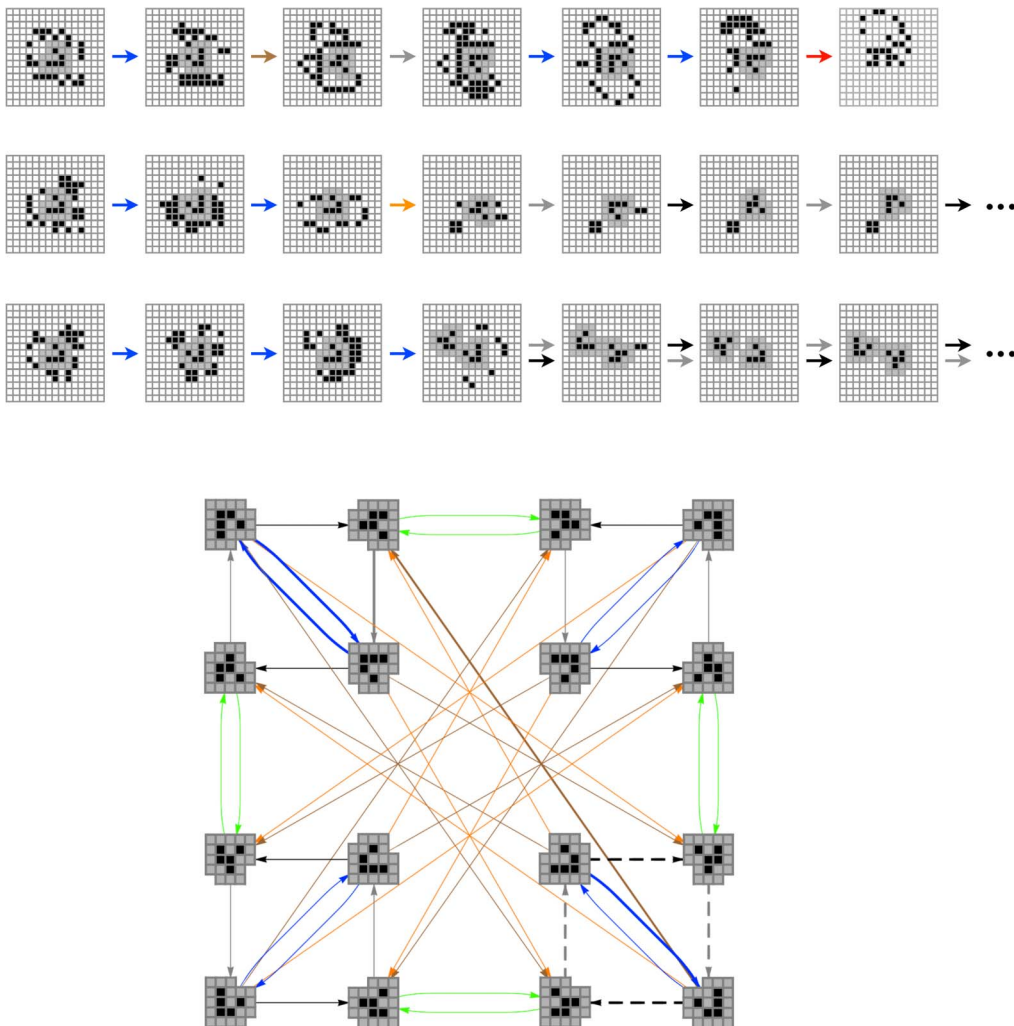
Figure 10. Three examples of structurally coupled interactions from a preliminary study of (1,2)-environments. The path through a glider's cognitive domain corresponding to the uppermost interaction is shown at the bottom, with dashed lines indicating the path the glider would have followed in isolation and thick lines indicating its structurally coupled path.
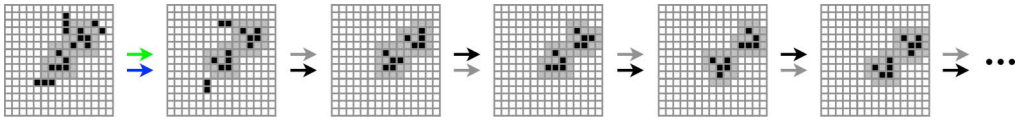
Figure 11. The simplest nontrivial example of a communicative interaction between two gliders supported by some surrounding environmental scaffolding.

{BLUE, BLUE, BLUE} sequence before spawning a second glider whose boundary initially overlaps with its own. Both gliders then enter into null perturbation loops, moving away from one other and persisting forever.

Even this preliminary study of the trajectories of structural coupling has revealed some intriguing general patterns. First, the trajectories that survive to a given iteration are always a tiny fraction of the possible perturbation sequences. In addition, of the trajectories that do survive to a given iteration, many fall into similar sequences after initial transients have passed. For example, a glider eventually leaving behind a single block, as shown in the middle example in Figure 10, occurs frequently. Furthermore, some sequences of nonnull perturbation classes seem to occur much more frequently than others. For example, sequences of BLUE perturbations are quite common. Thus, although the combinatorics are daunting, it is clear that a more complete study of the trajectories of structural coupling might yield interesting insights into the structure of the set of all possible lives that a glider can live in a periodic lattice of a given size.

8 Communicative Interactions and Consensual Domains

Once we begin to consider larger environments, it becomes possible for one glider to encounter another. Indeed, we have already seen an example of this in the bottom sequence in Figure 10. Recall that, for MV, entity-entity perturbations are communicative interactions that serve to orient entities within their respective cognitive domains to new possibilities for action. When these communicative interactions are mutual and ongoing, the participating agents generate a new consensual domain of coordinated orientations. Can gliders in GoL be used to explore any of these ideas?

It is relatively easy to see that two gliders can only directly interact when their boundaries partially overlap, since it is only then that the ON cells of each glider intersect the 1-environment of the other. It is also straightforward to enumerate all such overlaps and examine their consequences. In many cases, one or both participating gliders disintegrate after one step. However, there are some instances in which both gliders survive the interaction, sometimes overlapping for multiple steps. Unfortunately, the interactions between two such gliders are always null perturbations (i.e., BLACK or GRAY); from the perspective of each glider it is as if the other one is not even present. In fact, this will always be the case, because, as pointed out in Section 5, nonnull perturbations require nonlocal patterns of ON cells and two gliders can perturb one another only locally. Thus, no nontrivial communicative interactions are possible between two gliders in an otherwise empty lattice.

Nonnull communicative interactions can occur in more complex environments containing either more than two gliders or additional environmental scaffolding. A simple example of the latter is shown in Figure 11, where a W_0^R and an R_0^R glider interact with one another and a nonempty environment. Here the central W_0^R glider experiences a BLUE perturbation, causing it to transition to state $W_{\pi/2}^L$ instead of R_π^R. At the same time, the R_0^R glider experiences a GREEN perturbation, causing it to transition to state R_π^L instead of $W_{\pi/2}^R$. From this point on, both gliders not only persist but also maintain an ongoing interaction through their overlapping boundaries, jointly moving toward the lower right corner of the lattice 90° out of phase. Thus, due to their initial communicative interaction, their joint states have fallen into a coordinated pattern of action that defines a consensual domain. Unfortunately, all communicative interactions beyond the first one correspond to null perturbations.

This brief exploration has demonstrated that nontrivial communicative interaction between gliders is possible, but much more work needs to be done to determine whether the chains of such nontrivial interactions required to generate an interesting consensual domain can occur. Ideally, we would like a configuration in which at least two gliders surrounded by appropriate environmental scaffolding could engage in a recurring sequence of nontrivial communicative interactions. Then we could characterize each glider's path through its cognitive domain as a consequence of its communicative interactions with the other glider. We could also examine the structure of the resulting consensual domain that their sequence of joint states defines. Unfortunately, the challenge here is the same as for structural coupling. We need to search the space of GoL configurations for those with desired properties in the face of daunting combinatorics. One interesting idea might be to try to adapt some of the glider-manipulation technology that GoL enthusiasts have developed over the years for changing the phase or direction of motion of a glider (www.conwaylife.com/wiki). However, in every case so far examined, this technology destroys the original glider and then creates another one with the desired properties several iterations later, violating the fundamental autopoietic constraint of maintaining an identity.

9   Discussion

The central goal of this article has been to illustrate in some technical detail how many of the key ideas underlying MV's biology of cognition play out in a simple concrete model. First, we presented a conception of a glider in the game of Life as a spatiotemporal pattern of ON and OFF cells that participate in a closed network of interactions that generate and maintain that very same pattern. This definition of a glider also specifies a boundary, which in turns defines an immediate environment that serves as an interface between the glider and the rest of the lattice. Second, by explicitly enumerating the consequences to a glider of every possible perturbation of its immediate environment, we completely described the structure of a glider's cognitive domain as the set of all nondestructive interactions into which it can enter. Third, we found that this cognitive domain divides the nondestructive perturbations into six classes, which form a glider's *Umwelt*, and that the structure of each perturbation class can be derived explicitly from the constraints that a glider must satisfy in order to maintain its existence. Fourth, we characterized the interlocking sequences of mutual perturbation that occur during structural coupling between a glider and its environment and demonstrated how the tools of information theory can be used to quantify the resulting chains of dependencies. Finally, we suggested how one might be able to explore consensual domains in this model by presenting a very simple example of a communicative interaction between two gliders.

One point to emphasize about our analysis is that the state change a glider undergoes at each point in time is codetermined by its own internal dynamics and the perturbations that it receives from its environment. Although the very limited number of states that a glider can be in may obscure this fact, an external perturbation alone cannot in general uniquely determine its next state. For example, the same BLUE perturbation that merely reorients a glider in a W state completely destroys a glider in an R state. Another important manifestation of this point is the significant imbalance in robustness to perturbation between W states (which can survive 72,440 distinct microperturbations) and R states (which can survive only 22,496).

An important issue that deserves a deeper exploration is how best to define the possible states of a glider. There are at least three possibilities. First, if we grant significance to the absolute coordinates of each lattice cell, then there are 16 possible gliders centered about every cell in the lattice, for a total of $16NM$ possible glider states in a periodic $N \times M$ lattice and an infinite number of possible glider states in an unbounded lattice. Second, if we grant significance only to relative positions within the lattice (i.e., we abstract over translational symmetry), then we obtain the 16 possible glider states that we have used in this article (Figure 3). The argument for this definition is that, although the absolute coordinate system we impose on the lattice would seem to be completely arbitrary, the orientation and chirality of a glider are relevant to how we as external observers describe how a

glider relates to its environment. Finally, if we also remove the significance of orientation and chirality from our definition of glider state (i.e., we also abstract over rotational and reflectional symmetries), then we are left with only two states: W and R. This is arguably the appropriate definition of a glider's intrinsic state, since, from an internal perspective, a glider has no way of determining its orientation or chirality. From this internal perspective, the six macro perturbation classes that we have explored in this article collapse to only four (BLUE, GREEN, GRAY, and BLACK/BROWN/ORANGE), and the glider's cognitive domain becomes much simpler. However, as mentioned above, the distinction between W and R forms still has important behavioral consequences.

The distinction we drew in Section 4 between the microdynamics of lattice cells (as defined by f) and the macrodynamics of glider states (as defined by F) nicely illustrates a point that has been emphasized previously [1]: If we are only interested in the behavioral dynamics of an agent whose existence we are willing to take for granted, then this behavioral dynamics can be characterized by the tools of dynamical systems theory. Furthermore, such a behavioral characterization can often be performed in relative isolation from an account of the constitutive processes that underlie this behavioral dynamics, as we have done here, for example, in Section 6 (although this may not always be the case [6]). This observation justifies the abstractions that are common in situated, embodied, and dynamical approaches to behavior and cognition. However, if we wish to understand why an agent's behavioral dynamics has the particular structure that it does or we wish to probe an agent's behavior at the limits of its viability, then we must consider at least some aspects of its underlying constitution.

More generally, how might the key ideas of the enactive approach to cognition manifest themselves in this simple model? At its core, enaction is about a self-sustaining entity bringing forth, through its history of structural coupling with its environment, a world of significance and a value-laden point of view on that world [26, 28, 31]. The analysis in this article has demonstrated how the closed network of processes that constitute a glider classify environmental configurations into non-destructive and destructive subsets that are respectively "good" and "bad" for the glider. In addition, we have seen how a glider's constitution further classifies the nondestructive perturbations into six subclasses according to which structural changes it undergoes in response. These categories do not preexist as objective properties of the GoL universe. Rather, they are induced by the existence of a glider. This brings a subjective or phenomenal aspect into even this simple model. A particularly key feature of this simple model of enaction is that it allows us to analyze an agent from both the "inside" and the "outside" simultaneously, making it possible to reconcile the subjective and objective perspectives in a way that we currently find extremely difficult to do with ourselves. Thus, although much work remains to be done, at least some of the key ideas of the enactive approach might be usefully explored in this simple model.

Interestingly, Di Paolo [5] has questioned the sufficiency of the sorts of autopoietic ideas that we have examined here for enaction. In particular, he argues that a merely autopoietic system is incapable of exhibiting the graded norms of better and worse; only the binary norm of survival versus disintegration is implied by autopoiesis. Di Paolo suggests that the concept of adaptivity—"the capacity of an organism to regulate itself with respect to the boundaries of its own viability"—must be explicitly added to autopoiesis in order to fully capture the sense-making required for true enaction. Although full justice cannot be done to this debate here, it is worth pointing out how careful consideration of simple models can illuminate otherwise subtle conceptual issues such as this.

It is certainly true that a glider's constitution fundamentally distinguishes between destructive and nondestructive perturbations. But it does much more than that. Within the class of nondestructive perturbations, a glider further distinguishes subclasses according to the state within its cognitive domain that results from each such perturbation. This is a simple consequence of structural degeneracy: The same organization can be realized by multiple structures. Of course, the structural degeneracy of a glider is quite mild, but this degeneracy can be astronomical in more complicated entities. These differences in state in turn have consequences for the glider's sensitivity to subsequent perturbations and thus implicitly constitute different valuations of its situation. A glider prospers in (i.e., maintains its structural coupling with) precisely those environments to whose

spatiotemporal structure its autopoietic dynamics is matched. A persisting autopoietic system therefore necessarily behaves as if it is changing its future responses as a consequence of its past interactions. The fact that these conclusions follow directly from MV's original autopoietic framework suggests that the necessity of adding an explicit notion or process of adaptivity to that framework may require further discussion. Indeed, this observation is nothing more than a reiteration of the central lesson of work on learning without synaptic plasticity: Dynamics alone is sufficient for adaptivity; no explicit adaptive process is required [12, 20, 32].

As simple as it is, the game of Life has turned out to provide a rich medium for exploring some of the key ideas underlying MV's biology of cognition and the enactive approach that has grown out of it. Of course, this simple model cannot do adequate justice to the full range of topics in this literature. At best, it can perhaps serve as a starting point for concretely grounding the conceptual analysis that lies at the foundations of this approach. Thus, there are many directions for future development. First, it would be very interesting to formalize the argument sketched in Section 3 that a glider is in fact a legitimate model of autopoiesis by mathematically characterizing the form of its organizational closure and showing that it satisfies MV's definition within the GoL lattice. Second, much more work remains to be done on describing the set of possible trajectories of structural coupling and the interlocking chains of dependence between glider and environment that they produce. Third, in order to further explore the notion of a consensual domain it will be necessary to find longer chains of nonnull communicative interactions between sets of gliders and their shared environment. Fourth, it would be worthwhile to apply the analysis in this article to other self-producing/maintaining patterns in the game of Life and other cellular automata in order to catalog the different possible cognitive domains and the trajectories of structural coupling in which they can engage. Finally, it will ultimately be important to extend this analysis to more complicated and realistic models of autopoiesis that would more completely capture the many facets of the enactive approach to cognition.

Acknowledgments

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

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

3. Bourgine, P., & Stewart, J. (2004). Autopoiesis and cognition. *Artificial Life*, *10*(3), 327–345.

4. Cover, T. M., & Thomas, J. A. (2006). *Elements of information theory* (2nd ed.). Hoboken, NJ: Wiley.

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

6. Egbert, M. D., & Barandiaran, X. E. (2011). Quantifying normative behaviour and precariousness in adaptive agency. In T. Lenaerts, M. Giacobini, H. Bersini, P. Bourgine, M. Dorigo, & R. Doursat (Eds.), *Proceedings of the 11th European Conference on Artificial Life* (pp. 210–218). Cambridge, MA: MIT Press.

7. Egbert, M. D., Barandiaran, X. E., & Di Paolo, E. A. (2010). A minimal model of metabolism based chemotaxis. *PLoS Computational Biology*, *6*(2), e1001004. doi:10.1371/journal.pcbi.1001004

8. Egbert, M. D., & Di Paolo, E. A. (2009). Integrating behavior and autopoiesis: An exploration in computational chemo-ethology. *Adaptive Behavior*, *17*(5), 387–401.

9. Fontana, W., & Buss, L. W. (1994). The arrival of the fittest: Toward a theory of biological organization. *Bulletin of Mathematical Biology*, *56*(1), 1–64.

10. Froese, T., & Ziemke, T. (2009). Enactive artificial intelligence: Investigating the systemic organization of life and mind. *Artificial Intelligence*, *173*(3–4), 466–500.

11. Gardner, M. (1983). *Wheels, Life and other mathematical amusements*. New York: W. H. Freeman.

12. Izquierdo, E., Harvey, I., & Beer, R. D. (2008). Associative learning on a continuum in evolved dynamical neural networks. *Adaptive Behavior*, *16*(6), 361–384.

13. Luisi, P. L. (2006). *The emergence of life: From chemical origins to synthetic biology*. Cambridge, UK: Cambridge University Press.

14. Maturana, H. R. (1978). Biology of language: The epistemology of reality. In G. A. Miller & E. Lenneberg (Eds.), *Psychology and biology of language and thought: Essays in honor of Eric Lenneberg* (pp. 27–63). New York: Academic Press.

15. Maturana, H. R., & Varela, F. J. (1980). *Autopoiesis and cognition*. Dordrecht, The Netherlands: D. Reidel.

16. Maturana, H. R., & Varela, F. J. (1987). *The tree of knowledge*. Boston: Shambhala.

17. Mavelli, F., & Ruiz-Mirazo, K. (2007). Stochastic simulations of minimal self-reproducing cellular systems. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *362*(1486), 1789–1802.

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

19. Ono, N., & Ikegami, T. (2000). Self-maintenance and self-reproduction in an abstract cell model. *Journal of Theoretical Biology*, *206*(2), 243–253.

20. Phattanasri, P., Chiel, H. J., & Beer, R. D. (2007). The dynamics of associative learning in evolved model circuits. *Adaptive Behavior*, *15*(4), 377–396.

21. Poundstone, W. (1984). *The recursive universe*. New York: William Morrow.

22. Stewart, S., Gapenne, O., & Di Paolo, E. (Eds.). (2010). *Enaction: Toward a new paradigm for cognitive science*. Cambridge, MA: MIT Press.

23. Suzuki, K., & Ikegami, T. (2009). Self-repairing and mobility of a simple cell model. In J. Pollack, M. Bedau, P. Husbands, T. Ikegami, & R. A. Watson (Eds.), *Artificial Life IX: Proceedings of the Ninth International Conference on the Simulation and Synthesis of Living Systems* (pp. 421–426). Cambridge, MA: MIT Press.

24. Thompson, E. (2007). *Mind in life*. Cambridge, MA: MIT Press.

25. Varela, F. J. (1979). *Principles of biological autonomy*. New York: North Holland.

26. Varela, F. J. (1988). Structural coupling and the origin of meaning in a simple cellular automaton. In E. E. Sercarz, F. Celada, N. A. Mitchison, & T. Tada (Eds.), *The semiotics of cellular communication in the immune system* (pp. 151–161). Berlin: Springer-Verlag.

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

28. Varela, F. J., Thompson, E., & Rosch, E. (1991). *The embodied mind*. Cambridge, MA: MIT Press.

29. Virgo, N. (2011). *Thermodynamics and the structure of living systems*. Unpublished doctoral dissertation, University of Sussex, Brighton, UK.

30. von Uexküll, J. (1934). A stroll through the worlds of animals and men: A picture book of invisible worlds. Reprinted in 1957 in C. L. Schiller (Transl. & Ed.), *Instinctive behavior: The development of a modern concept* (pp. 5–80). International Universities Press.

31. Weber, A., & Varela, F. J. (2002). Life after Kant: Natural purposes and the autopoietic foundation of biological individuality. *Phenomenology and the Cognitive Sciences*, *1*(2), 97–125.

32. Yamauchi, B., & Beer, R. D. (1994). Sequential behavior and learning in evolved dynamical neural networks. *Adaptive Behavior*, *2*(3), 219–246.

33. Zeleny, M. (1977). Self-organization of living systems: A formal model of autopoiesis. *International Journal of General Systems*, *4*(1), 13–28.