

Provenance of life: Chemical autonomous agents surviving through associative learningStuart Bartlett ^{*}*Division of Geological and Planetary Sciences, California Institute of Technology, Pasadena, California 91125, USA
and Earth-Life Science Institute, Tokyo Institute of Technology, Tokyo 152-8550, Japan*David Louapre *Ubisoft Entertainment, 94160 Saint-Mandé, France and Science Étonnante, 75014 Paris, France[†]*

(Received 31 March 2022; accepted 21 June 2022; published 12 September 2022)

We present a benchmark study of autonomous, chemical agents exhibiting associative learning of an environmental feature. Associative learning systems have been widely studied in cognitive science and artificial intelligence but are most commonly implemented in highly complex or carefully engineered systems, such as animal brains, artificial neural networks, DNA computing systems, and gene regulatory networks, among others. The ability to encode environmental information and use it to make simple predictions is a benchmark of biological resilience and underpins a plethora of adaptive responses in the living hierarchy, spanning prey animal species anticipating the arrival of predators to epigenetic systems in microorganisms learning environmental correlations. Given the ubiquitous and essential presence of learning behaviors in the biosphere, we aimed to explore whether simple, nonliving dissipative structures could also exhibit associative learning. Inspired by previous modeling of associative learning in chemical networks, we simulated simple systems composed of long- and short-term memory chemical species that could encode the presence or absence of temporal correlations between two external species. The ability to learn this association was implemented in Gray-Scott reaction-diffusion spots, emergent chemical patterns that exhibit self-replication and homeostasis. With the novel ability of associative learning, we demonstrate that simple chemical patterns can exhibit a broad repertoire of lifelike behavior, paving the way for *in vitro* studies of autonomous chemical learning systems, with potential relevance to artificial life, origins of life, and systems chemistry. The experimental realization of these learning behaviors in protocell or coacervate systems could advance a new research direction in astrobiology, since our system significantly reduces the lower bound on the required complexity for autonomous chemical learning.

DOI: [10.1103/PhysRevE.106.034401](https://doi.org/10.1103/PhysRevE.106.034401)**I. INTRODUCTION AND CONCEPTUAL FRAMEWORK**

Ever since the outspoken chemist Ilya Prigogine introduced the concept, dissipative structures and their underlying phenomenology have beguiled, entranced, and split opinions of researchers from a range of disciplines [1,2]. The approximate definition of a dissipative structure is a coherent and discernible dynamical pattern that is maintained by external disequilibria over a finite time. The structure exists at a higher phenomenological or descriptive level to its microscopic constituents and can thus be considered “emergent” [3–7].

The most frequently cited examples of dissipative structures are fluid convection cells [8–12], hurricanes [13], dynamic surfactant structures such as micelles, vesicles, and droplets [e.g., Refs. [14–19]], stars and galaxies [20], black holes [21], and biological organisms [e.g., Refs. [22–24]]. In the fields of artificial life and the origins of life (OoL), dissipative structures are a key focus area, serving as metaphors for simple forms of life and providing clues to the nonlife to life transition [15,25–28]. Numerous experiments and simu-

lations in this field have illustrated the emergence of subsets of life’s properties. For example, various oil droplet systems have been shown to exhibit chemotaxis, in which the motion of the droplets is powered through the consumption of a fuel compound, and the droplets’ motion follows gradients in the concentration of that fuel [15,16,29–33]. Self-replication is a common phenomenon in reaction-diffusion systems (RDSs) [24,27,34–40] and vesicle-based structures, also known as protocells [41–50]. Nonliving or artificial dissipative structures also exhibit ecological behavior including competition [34,36,51], homeostasis [52] and symbiosis [35,53]. Artificial cell studies commonly follow the doctrine of autopoiesis [54], which is closely related to the “chemoton” concept [55] and the container-metabolism-program framework [56], in that they seek integrated cellular structures comprising a boundary or membrane, a metabolic system for converting precursor ‘food’ molecules into the components of the cell, and an information system that is normally analogous to a genetic system [45,46,50,57–60]. While this approach is philosophically sound, the simultaneous synthesis of all the required chemical components for such systems has proven to be prohibitively difficult, especially in “prebiotically plausible” scenarios, normally requiring very high precursor concentrations [61,62]. Additionally, the physical driving forces that

^{*}sjbart@caltech.edu[†]<https://scienceetonnante.com>

would compel a system of molecular components to exhibit the requisite dynamics that we associate with life are poorly understood [63,64].

While pursuing a parsimonious, “straight-shot” from prebiotic physics and chemistry to biology is a logical starting point in seeking the OoL, given the often bizarre twists and turns that life’s evolution has taken [65–67], it is also possible that the earliest forms of life were not simplified versions of extant life, but rather different in composition and organization [68,69]. A natural question thus arises: What are the conserved quantities that we expect extant life, early life, and even extraterrestrial life, to exhibit? While definitions of life abound [70–76], and much previous work was inspired by the autopoiesis and chemoton concepts, a definition has recently been introduced that combines traditional and modern ideas from thermodynamics, biology, information theory and cognitive science [77]. This ‘four pillared’ definition suggests that a living system must exhibit all of the following properties: (1) dissipation (the system must be exposed to one or more thermodynamic disequilibria or free energy sources), (2) autocatalysis (the system must exhibit or have the capacity to exhibit exponential growth of a representative size metric under ideal conditions), (3) homeostasis (the system must possess regulatory or negative feedback mechanisms that can mitigate external or internal perturbations), (4) learning (the system must have the ability to sense, store, process and exploit information).

While various model protocell systems have primitive genetic components, and simpler artificial chemical systems exhibit self-replication, chemotaxis, and homeostasis, emergent learning in simple artificial systems has remained an elusive goal. Chemical computing itself is a large discipline that includes the incredible achievements of DNA computing [78–82], computing using the Belousov-Zhabotinsky reaction [83–91], and more general approaches that exploit the computational universality of chemical reaction networks [92–105]. Chemical computation has also been exploited for the control of agentlike entities such as simple robots [106], but such robots are not emergent and hence their analysis is primarily a heuristic tool for the engineering of swarm intelligence (as opposed to understanding the OoL).

Despite the great strides described above, the problem of the emergence of a self-organized entity that satisfies a minimal definition of life including a basic learning ability, is still very much open. Complex protocells that are scaled-down versions of extant life provide extensive but somewhat limited guidance to the OoL, since there remains a gap between the passive, blind worlds of physics and chemistry, and the active and goal-directed world of biology. In the present work, we explored a minimal complexity, emergent system that exhibits all the features of minimal life definitions, including the four pillars of “lyfe” [77], autopoiesis [54], and the ability to perform associative learning.

This result reduces the lower bound on the required system complexity of an autonomous chemical agent capable of learning. Our approach is also expandable to larger sets of variables or features. Given the vast diversity of prebiotic reaction networks, it is almost guaranteed that the required reactions can be found in a natural system, in the context of the OoL. The key challenge is understanding the conditions

in which learning becomes the most stable dynamical process in a system (the learning has to feed back positively on the stability of the patterns performing the learning, it has to be dynamically favored).

Our approach is based upon the solitonlike, spatially organized instabilities that form in nonlinear chemical systems, collectively known as reaction-diffusion structures (RDSTs) [24,27,34–39,53,107–122]. The systematic understanding of RDSTs traces all the way back to [123], who introduced their founding principles alongside the foundations of morphogenesis, which became one of the pioneering triumphs of mathematical biology. The RDS used in the present work is commonly known as the Gray-Scott model (GSM) [37,38,40,113,114,117,118], which is an elaboration of the Selkov model of glycolysis [124]. Experimental realizations of this system include the ferrocyanide-iodate-sulfite reaction [125–127], and the rich dynamics of heterogeneous catalysis on surfaces [128–133].

The GSM is a simple two-dimensional (2D) RDS that has been shown to exhibit several emergent, lifelike properties. It involves two chemical species, denoted A and B , interacting through a simple autocatalytic reaction $A + 2B \rightarrow 3B$. Along with an appropriate supply mechanism for A , and removal mechanism for B , this reaction produces a variety of patterns, including self-replicating spots, able to progressively expand into regions of space [40]. In addition, previous work has shown that adding the thermal dimension to this system (original versions were isothermal) reveals even more layers of emergent phenomena including competition (between RDSTs and convection cells) [36], and thermal homeostasis through symbiotic thermal regulation [35].

II. THE ROLE OF LEARNING IN THE BIOSPHERE AND ITS ORIGINS

Although the GSM exhibits lifelike characteristics, it lacks an essential feature of life: learning. Broadly speaking, learning can be defined as the ability of a system to record information about its environment and process that information to modulate its behavior. It is believed that learning is an essential pillar of life, providing a sensitivity to the environment that can improve the survival probability of the living system (this includes fields such as cognitive science, 4E cognition (embodied, embedded, enactive, and extended), perception-action cycles, measurement-feedback protocols, immune cognition, etc.) [134–150]. Although learning is sometimes viewed as specific to organisms having a nervous system, a large and rapidly expanding list of nonneural learning mechanisms illustrate its biological ubiquity, and it is hypothesized that learning may have played a role in life’s origins [151–155]. Among the different learning mechanisms, associative learning is particularly relevant in this regard. It can be defined as the ability of a system to detect and record *correlated* features about its environment [156–158]. Note that such correlations can be either positive (high levels of variable M are associated with high levels of variable N), or negative (high levels of M are associated with low values of N , and vice versa).

The selective advantage provided by such learning mechanisms can be illustrated by the case of predator-prey

interactions. The simplest strategy for a prey to escape a predator would be to trigger a response at the moment of the actual predator attack. This purely reactive behavior could be called a *direct* strategy. However, since predators are often relatively fast and the motor abilities of prey relatively limited, there is a survival incentive for prey species to be able to anticipate and react in advance.

To do so, prey species use their sensory systems to detect warning signs of predators. For instance, gazelle detect stalking predators such as cheetahs by visual detection, fish can sense sudden variations in light or fluid flow, etc. Detecting a predator ahead of time can then trigger a rapid response in order to escape. Such a *preemptive* strategy is clearly advantageous since it allows more time to escape, as opposed to waiting for the actual predator attack. However, since such an escape manoeuvre is energetically costly, prey species have an incentive to escape only when there is an actual threat, and avoid false positive detections.

The simplest way to achieve an escape response would be to react when a given sensory stimulus (for instance level of sound, or fluid flow disturbance) exceeds a certain threshold. However, that threshold will depend on the environment (e.g., background level of noise, water turbulence, etc.). As those environmental conditions may slowly vary because of natural cycles or prey displacement, the threshold for reaction should be continuously adjusted. For instance, for a fish, a change of visibility caused by depth, turbidity or seasonal cycles should have an impact on the decision to trigger a fast escape when a given visual stimulus is sensed. Achieving such a behavior as the environment changes is an *associative learning* task: prey species need to learn the extent to which a stimulus is associated with an actual threat in their current environmental conditions, and update their behavior if environmental conditions happen to change.

In this framework of predator-prey interactions, we consider three possible strategies: a direct reaction, a preemptive reaction, and associative learning. Although the benefit of such an associative learning ability in the case of a predator-prey interaction is quite clear, let us illustrate how this could also play a role in the resilience and adaptability of a protoliving chemical system.

Consider a system that can be damaged by a chemical species T (a toxin), which is delivered at regular intervals. If the system has the ability to produce an antidote N , that degrades the toxin, then the simplest defense strategy would be the *direct* mechanism, where antidote N is directly produced when toxin T is encountered, for instance, through a catalytic reaction:



Note that throughout this work, it is assumed that mass conservation is satisfied through additional species that are omitted from the reaction expressions for simplicity (in this case, a precursor from which N is synthesized). Although such a direct network would be able to tune its response to the strength of the attack, it does not involve any anticipation. If the toxin is delivered in a sudden way at high concentration, then antidote production might be too slow to avoid significant damage to the system.

Now suppose that the delivery of toxin T is always preceded by the delivery of another chemical compound S . The presence of S could serve as a stimulus signaling the occurrence of T ahead of time. A possible *preemptive* strategy would then be to start producing antidote N as soon as stimulus S is detected, so the antidote concentration has time to build up in advance. This preemptive mechanism is equivalent to triggering a fast start escape whenever a predator is detected, rather than waiting for the actual predator strike. In our chemical system, such a preemptive mechanism can be realized by a reaction that produces antidote N and is catalyzed by stimulus S :



Now there is an obvious drawback with this strategy: since antidote production is controlled by the presence of stimulus, it is not sensitive to the actual strength of the attack (in our case, the concentration of toxin T). Suppose that the relative concentrations of S and T depend on slowly varying environmental conditions (like temperature, pH, or the concentration of another compound) such that a given stimulus concentration would not always correspond to the same amount of toxin. We could even imagine that under certain conditions, the association disappears so that stimulus remains but toxin occurrences cease. A simple preemptive network, relying on stimulus concentration alone, would produce antidote irrespective of the actual toxin level, and would even keep producing antidote after the association between S and T disappeared. This unnecessary production of antidote would be detrimental, as it would incur a metabolic free energy cost, or lead to damage if the antidote itself is a mild toxin. This type of spurious antidote production would be analogous to a prey animal repeatedly triggering fast start escapes even in the absence of a real threat. Although many prey species of mammal do exhibit this type of nervousness (even humans can be easily startled), they are clearly not fleeing the scene at every hint of a predator's presence.

To avoid this, a more efficient chemical network would have the ability to detect and record whether, in the prevailing environmental conditions, S is actually followed by T , and to what extent (the strength of the association). Such a network would then adapt if the environmental conditions were to change, for instance if the relative concentrations of S and T are evolving, or if the delivery of T is suppressed while that of S remains. Fulfilling such a requirement is an associative learning task: the network should first detect to what extent S is followed by T , record that information and react accordingly, while being able to update itself whenever the link between S and T changes. We can hypothesize that for a large class of conditions, such an *associative learning* network would confer greater resilience than any direct or preemptive network, providing a selective advantage for emergent systems capable of associative learning.

As described above, chemical computing is already a well-developed field, and approaches to chemical learning have been suggested [98–101,104]. However, the approach used in the present work was inspired by the learning networks of Ref. [159]. Using search methods inspired by evolutionary processes, these authors were able to find abstract chemical networks able to solve learning tasks (such as basic

association, and the AB-BA task). However, their model used a very large parameter space (searched using an evolutionary algorithm), and it was somewhat difficult to interpret how to choose the values of the different parameters involved. Furthermore, the learning networks were not embodied within emergent chemical structures.

The present work aims to: (1) explore minimally complex reaction systems that clearly exhibit associative learning, (2) investigate the resilience of such networks compared to simpler strategies, such as the direct or preemptive networks described above, (3) assess whether such networks can be embedded within emergent, spatially distinct, self-organised chemical structures. Such phenomena would serve as an important guide to the conditions in which the basic properties of life can emerge. Crucially, this emergence occurs in the absence of complex molecules, peptides, nucleic acids, biochemical reactions, metabolic cycles, or membrane structures, and hence lowers the bound on the necessary conditions for chemical, autonomous, cognitive agents.

In the following section, we present minimal chemical networks that efficiently solve the learning task described above, when placed in a well-stirred [zero-dimensional (0D)] environment. We display plausible external conditions in which those networks outperform direct or preemptive strategies. In the second part of the paper, we extend those results to the two-dimensional case, where the learning network is coupled to a self-replicating spot system from the GSM. This provides a realization of a dissipative structure able to replicate and adapt to its environment using associative learning, a feat which, as far as we are aware, has not been achieved before.

III. CHEMICAL LEARNING NETWORKS IN WELL-STIRRED, ZERO-DIMENSIONAL SYSTEMS

A. Environment and learning task

Let us first describe the task to be solved. Consider an environment where a stimulus S and a toxin T are regularly delivered with period \mathcal{T} , and a temporal separation $\Delta\mathcal{T}$ between them, such that S can be used as a signaling cue for the later occurrence of T (we choose $\mathcal{T} = 1000$ and $\Delta\mathcal{T} = 400$). We consider that S is delivered as boluses of unit concentration, while T is delivered as boluses of concentration ϵ . Finally, we assume both S and T degrade exponentially with a short characteristic time ($\tau_S = 0.2\mathcal{T}$), so that they quickly disappear after each bolus. To model slowly varying environmental conditions that manifest as a modulation of the association between stimulus and toxin, the concentration ϵ of toxin boluses is allowed to slowly vary with time. We consider $\epsilon(t)$ linearly rising from 0 to 1 during the first 10 delivery periods, decreasing linearly during the next 10 periods, and staying at zero thereafter. This is illustrated in Fig. 1. We now consider a chemical compound B , serving as a proxy for the presence of a protoliving chemical system. To achieve this, B has an autocatalytic dynamic, with an equilibrium concentration of unity, representing environmental capacity constraints. Evolution of the concentration of B , denoted ψ_B , is modeled by a logistic growth differential equation:

$$\frac{d\psi_B}{dt} = \frac{1}{\tau_B} \psi_B (1 - \psi_B). \quad (3)$$

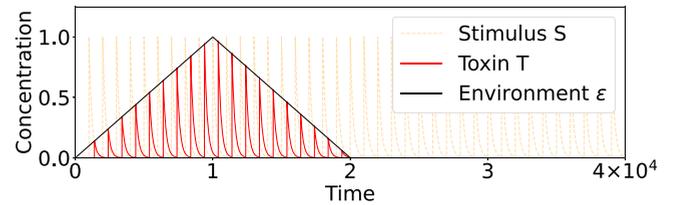


FIG. 1. Dynamics of key compounds in learning task. Stimulus and toxin bolus deliveries in the learning task to be solved. Delivery period is $\mathcal{T} = 1000$ and the time separation between stimulus and toxin occurrences is $\Delta\mathcal{T} = 400$. Stimulus boluses are always of unit concentration, while the magnitude of toxin boluses are controlled by the environmental parameter $\epsilon(t)$ (black line), which varies slowly over time. Both stimulus S and toxin T decay exponentially with a short characteristic time $\tau_S = 0.2\mathcal{T}$.

Note that the logistic differential equation is a realization of the following set of chemical equations: $B \longrightarrow B + B$ (autocatalysis) and $B + B \longrightarrow \emptyset$ (capacity constraint).

The growth rate of ψ_B vanishes as it reaches its equilibrium value of 1, and τ_B represents the (minimum) characteristic growth time. We choose $\tau_B = 500$, so that the system has time to recover between each toxin delivery, but not completely.

To model the effect of toxin on our system, T reacts with B , causing decay into inert waste products:



where the reaction constant $k_{BT} = 0.5$. This moderate value ensures that the protoliving system is significantly damaged by the toxin at each delivery, without being completely destroyed by a single event.

The evolution of B in the presence of stimulus and toxin delivery is illustrated in Fig. 2. We see that the system gets progressively destroyed after each toxin bolus. Despite partial recovery between each delivery (due to the autocatalytic behavior of B), the system (represented by the concentration of B) does not survive more than a couple of periods. Suppose now that T can be degraded by an antidote N , via a reaction into inert waste:



where $k_{NT} = 1$ so that the antidote is effective at removing T . To allow for functional anticipation mechanisms, we assume that the antidote is a long-lasting compound, and hence decays with a relatively long characteristic time $\tau_N = 5000$.

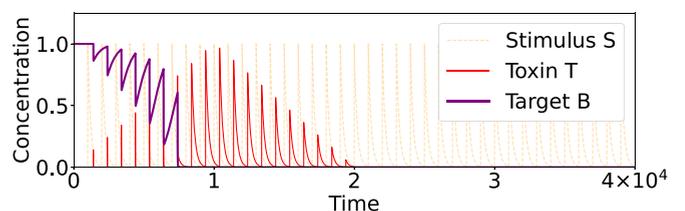


FIG. 2. Dynamics of system in the absence of a defense mechanism. Time evolution of the concentration of B with toxin deliveries, in the absence of any defense network (network \mathcal{N}_0 , see Table 1 for a summary of chemical equations used).

To implement the assumption that spurious antidote production is detrimental, we add a reaction in which the antidote degrades B as well,



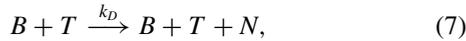
although with a much smaller rate constant $k_{NB} = 0.0025$. There are two interpretations of this effect: (1) the antidote itself acts as a mild toxin, or (2) production of the antidote incurs a metabolic (free energy) cost to the prebiotic system, hence it consumes a certain quantity of B via the above reaction. This will enforce the fact that antidote cannot be freely and arbitrarily produced, and must be used in a parsimonious and timely manner.

B. Reaction network design

Several defense networks can be considered for the production of antidote in response to the toxin. Below we describe a *direct* network \mathcal{N}_D , a *preemptive* network \mathcal{N}_P , and an *associative learning* network \mathcal{N}_A

1. Direct network

The simple, “instantaneous” response of network \mathcal{N}_D only involves one reaction in which the production of antidote N is directly catalyzed by B in the presence of toxin T (the relevant supply compound for N is omitted):



where reaction constant k_D controls the rate of the response.

2. Preemptive network

This network \mathcal{N}_P , resembles the direct network, but the production of antidote N is instead catalyzed by the presence of stimulus S :



where k_P controls the rate of the response.

3. Associative network

This network \mathcal{N}_A , involves two additional species: M , acting as a short-term memory, and L , acting as a long-term memory. Production of M is catalyzed by the presence of stimulus S :



To serve as a short-term memory, its characteristic decay time should be of the order of the expected separation time between the boluses of S and T , hence $\tau_M = 200$.

The long-term memory L will encode the association between stimulus and toxin. To do so, L will be produced from M in the presence of T (effectively acting as an AND gate on the simultaneous presence of M and T), hence recording the fact that a bolus of S has actually been followed by the occurrence of T :



Note that for the simple environment used here, it is possible to implement an associative network in which long-term

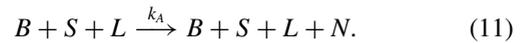
TABLE I. List of chemical equations used in each network.

Network		Equation numbers
\mathcal{N}_0	(No network)	(3) (4) (5) (6)
\mathcal{N}_D	(Direct network)	(3) (4) (5) (6)+(7)
\mathcal{N}_P	(Preemptive network)	(3) (4) (5) (6)+(8)
\mathcal{N}_A	(Associative network)	(3) (4) (5) (6)+(9) (10) (11)

memory is stimulated only by the presence of T , precluding the need for short-term memory M . This network is explored in the Supplemental Material [160], which also discusses why we retain short-term memory in network \mathcal{N}_A .

To serve as a long-term memory, the characteristic decay time of L should be of the order of several multiples of the period \mathcal{T} , reflecting the rate at which the system will learn and adapt to a varying environment. Hence, a value of $\tau_L = 4\mathcal{T} = 4000$ was chosen.

Finally, for this network production of antidote N is catalyzed by B and the long-term memory compound L , in the presence of stimulus S :



In this description of the associative network \mathcal{N}_A , three reaction rates k_M , k_L and k_A have been so far left unspecified, and can in principle be tuned to vary the dynamics and strength of the response. However, from the three reactions involved in the network, it can be seen that the amount of N produced essentially depends only on the product $k_M k_L k_A$. Hence, we choose to use only k_A as a free parameter when performing sensitivity analyses. Without loss of generality, we set $k_M = 0.01$ and $k_L = 0.25$ (those values being chosen so that maximum concentrations of M and L remain of order unity).

The chemical equations used in each network are summarized in Table I.

4. Simulation

We simulate the dynamics of the networks and the environment with a set of coupled ordinary differential equations. If $\{C_i\}_{i=1..N}$ represents the set of N chemical species, subject to R reactions, then the time evolution of concentration ψ_{C_i} is given by

$$\frac{\partial \psi_{C_i}}{\partial t} = - \sum_{r=1}^R k_r [\alpha'_{ri} - \alpha^*_{ri}] \prod_{j=1}^N \psi_{C_j}^{\alpha'_{rj}},$$

where α'_{ri} is the left-hand side stoichiometric coefficient for reaction r and chemical species i (the number of molecules of species i entering as reactants into reaction r), and α^*_{ri} is the right-hand side stoichiometric coefficient for reaction r and chemical species i (the number of molecules of species i leaving as products from reaction r). Here we have ignored the role of temperature, and k_r denotes the rate constant of reaction r . Reactions are simulated only in the forward direction, and possible reverse reactions are considered as extra, separate reactions from their forward counterparts. The exponential decay of each species is included as a reaction.

We simulated the networks coupled to the environment over 40000 time steps, starting with zero concentrations for

all species except unit concentration for B . We use explicit time forward Euler integration ($dt = 1$).

5. Optimization

For each network, only one reaction rate is left unspecified, and can be tuned to vary the strength of the response: k_D for the direct network \mathcal{N}_D , k_P for the preemptive network \mathcal{N}_P , and k_A for the associative network \mathcal{N}_A . Since our goal is to compare each of these networks at their maximum functionality, we searched for optimal values using line search optimization, with the temporal averaged concentration of B as the target to be maximised. Since T degrades B , such an optimization selects for values of reaction rates where the toxin causes minimal damage to the system. However, since B is slightly damaged by excess levels of antidote N , an optimal solution will also avoid spurious production of N .

C. Results

After appropriately tuning the reaction constants, we obtained the results displayed in Fig. 3, with $k_D = 0.1$ for \mathcal{N}_D , $k_P = 0.005$ for \mathcal{N}_P , and $k_A = 0.006$ for \mathcal{N}_A . We see that the direct network gets degraded after a few periods, and barely survives longer than when no network is present. The network is constrained by the production of antidote, and the toxin-induced decay of B (which is required to catalyze antidote production). Loss of B is thus too rapid to permit sufficient production of N . Note that there is an upper bound on the allowed rate constant k_D (see discussion at the end of this section).

In contrast, the preemptive network starts to build up antidote concentration as soon as stimulus boluses occur. This leads to the accumulation of antidote, which is then at a sufficiently high concentration when toxin strikes, allowing an effective removal of T , and limited degradation of B . However, the amount of antidote produced is only linked to stimulus concentration and does not depend on toxin level. This leads to an excess of antidote when the toxin boluses are small, and to spurious antidote production in the second half of the sequence, when stimulus remains but toxin is no longer delivered. This creates a high residual concentration of antidote, that continuously damages the system, constraining the concentration of B to a level significantly lower than its nominal equilibrium concentration $\psi_B = 1$.

Finally, we see that the associative learning network is more resilient than both the direct and the preemptive networks. As the concentration of toxin boluses increases over time, long-term memory L accumulates and enables the anticipatory production of antidote N when stimulus occurs. Furthermore, as the strength of the association between stimulus and toxin evolves, the long-term memory concentration changes accordingly and promotes a response of appropriate magnitude. As the association disappears, long-term memory decays to zero, as does antidote concentration, allowing a full recovery of the system toward its dynamic equilibrium. This provides an effective associative learning mechanism, recording information about the environment, using this information to react accordingly, while adapting behavior when the external condition changes.

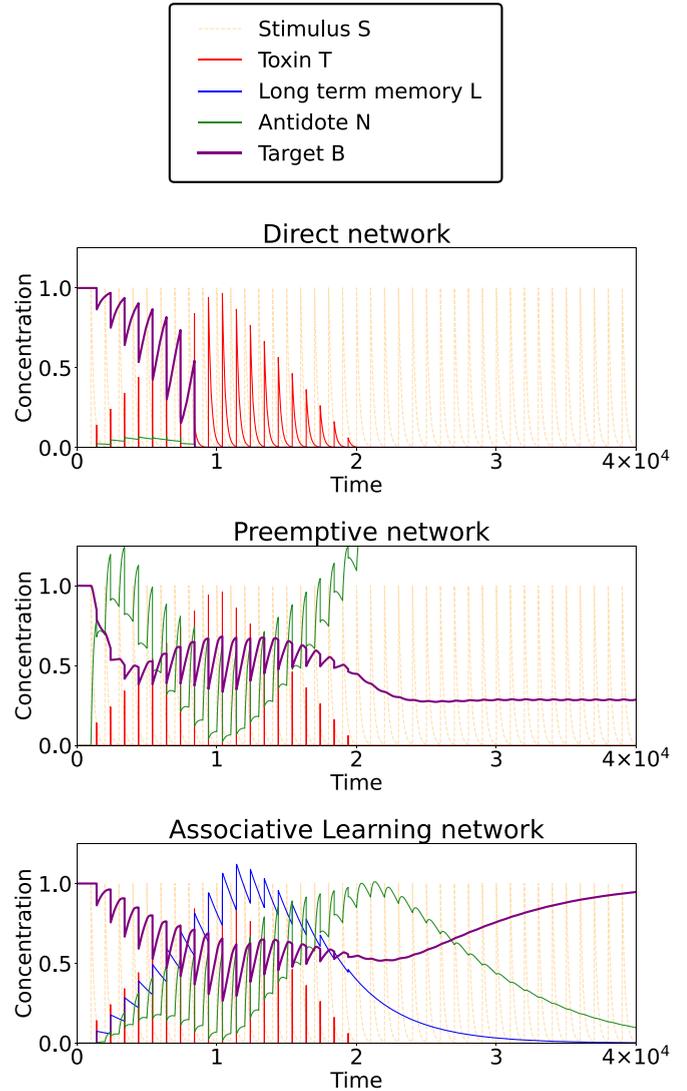


FIG. 3. Concentration time series for the three networks. Top: Direct network \mathcal{N}_D with $k_D = 0.1$. Middle: Preemptive network \mathcal{N}_P with $k_P = 0.005$. Bottom: Associative learning network \mathcal{N}_A with $k_A = 0.006$ (see Table I for a summary of chemical equations used in each network).

We can see that despite the learning mechanism, complete elimination of the toxin does not occur. This is because the system reaches a dynamic steady state through a feedback mechanism. Since the renewal of long-term memory is catalyzed by the presence of toxin, an excessive production of antidote would remove more of the toxin, and hence reduce the renewal of long-term memory, which would subsequently decrease the antidote level. Conversely, if the initial antidote level is too low due to excessive toxin, L would increase in response, with a concomitant increase in N . At steady state, the residual level of toxin provides just the necessary exposure to renew the long-term memory.

Moreover, variations in long-term memory concentration closely mimic the variation of the environmental parameter $\epsilon(t)$ (that drives toxin bolus concentration), and displays a high temporal correlation with it, although with a slight lag,

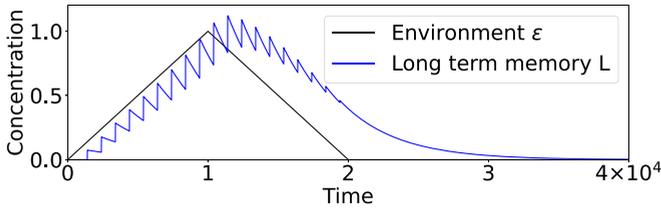


FIG. 4. Dynamics of environmental parameter and long-term memory. Time evolution of the environmental parameter $\epsilon(t)$ controlling toxin bolus concentration and long-term memory L for the optimal associative learning network.

as can be seen in Fig. 4. It shows that the associative network can effectively learn the value of $\epsilon(t)$ in a continuous fashion and adapt the system's behavior, while periodically updating its knowledge about the environment. Since the network revises its memory state (the set of concentrations of network compounds could be considered the system's "composome" [28,161]) at every bolus (with period \mathcal{T}), and has a natural response time τ_L , the evolution of L over time can be seen as the result of a recursive, discrete, low-pass filter applied to the (time-varying) magnitude of toxin boluses $\epsilon(t)$. This is a natural interpretation since at each bolus, the network samples ϵ as the strength of the association between S and T , and updates the level of L accordingly. This is exactly the function of a discrete low-pass filter, since the discrete low-pass filter of a signal x_n is a signal y_n such that $y_{n+1} = (1 - \alpha)y_n + \alpha x_n$. The parameter α governs the learning speed. For a low-pass filter with characteristic time $RC = \tau$ and discrete signal sampled with time step \mathcal{T} , we find $\alpha = \mathcal{T}/(\mathcal{T} + \tau)$. This can be illustrated further with the case of an environment parameter $\epsilon(t)$ that varies with a more complex temporal pattern, as shown in Fig. 5. It is instructive to consider the performance of each network when the reaction constant is varied. This sensitivity analysis is summarised in Fig. 6. When the reaction constant is too low, all three networks perform poorly, since the system gets quickly destroyed by toxin. The direct network performs poorly for all values of the reaction constant k_D . When its reaction constant k_P is sufficiently high, the preemptive network is able to neutralise the toxin. However, at higher values of k_P , its performance degrades because of damage caused by spurious antidote production at low or zero toxin level.

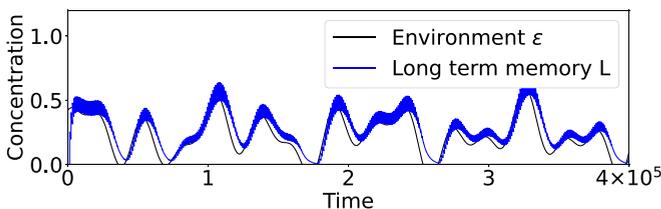


FIG. 5. Long-term memory concentration as a low-pass filter. Environmental parameter $\epsilon(t)$ varies with a more complex time evolution (computed as the sum of three sine waves of random amplitude, frequency and phase). The long-term memory concentration acts like a low-pass filter on $\epsilon(t)$. It approximates $\epsilon(t)$ with a slight lag related to the long-term memory characteristic timescale, and also exhibits high frequency oscillations (appearing as the thick segment of the curve) due to the periodic bolus deliveries.

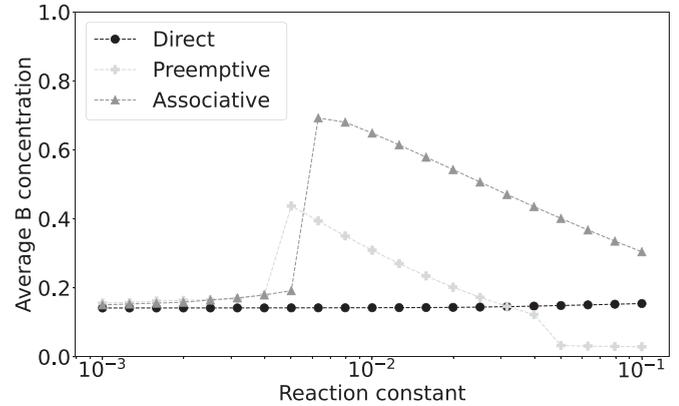


FIG. 6. Performance variations as a function of reaction constant. Performance of the different networks, measured as the temporal average of the concentration of B , as each key network reaction constant is varied. (See Table I for a summary of chemical equations used in each network.)

Finally, the associative network is able to perform effectively for a larger range of reaction constant values, especially since it does not suffer from spurious production of antidote when the association between stimulus and toxin disappears. This illustrates that the associative network outperforms both the preemptive network and the direct network in a wide range of situations. It can be seen that the performance of the direct network increases slightly with reaction constant, which is expected since a higher reaction constant allows that network to produce antidote at a higher rate when the toxin strikes. Increasing this reaction constant to even higher values would eventually allow that network to perform better. However, this would require a very high production rate of the antidote, typically a hundred times greater than that of the preemptive and associative networks. Reaching such extreme production rates is unlikely in realistic settings, where such synthesis would be limited by various factors such as precursor availability or saturation effects. If arbitrarily high reaction constants were allowed, then a direct network could eventually always perform as well as a preemptive or associative one. In the context of our predator-prey analogy, this would be equivalent to assuming that a prey animal can react and escape arbitrarily fast, in which case there is no need for it to learn to detect predators in advance. In practice, the speed at which a prey species can react is limited by environmental, physical, and biochemical constraints, hence the reason that associative learning became a widespread and important trait in the biological world, providing a competitive advantage over simple reactive strategies.

IV. CHEMICAL LEARNING NETWORKS IN 2D REACTION-DIFFUSION SYSTEMS

We now consider similar networks but embed them in a 2D environment, coupled to a GSM. Hence, there is a spatial dependence for the concentration of every species.

A. Model description

The model uses a 256×128 grid with a spacing of $dx = 1$, and an isotropic diffusion coefficient for each species.

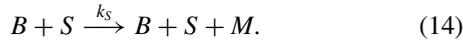
The GSM consists of two chemical species A and B, subject to the chemical reaction $A + 2B \longrightarrow 3B$, along with a decay mechanism for B and a supply mechanism for A [37,38,40,113,114,118]. The equations of motion for the concentrations of A and B are as follows:

$$\dot{\psi}_A = D_A \nabla^2 \psi_A - \psi_A \psi_B^2 + f(1 - \psi_A), \quad (12)$$

$$\dot{\psi}_B = D_B \nabla^2 \psi_B + \psi_A \psi_B^2 + (f + r)\psi_B, \quad (13)$$

with $D_A = 0.2$, $D_B = 0.1$, $f = 0.03$, and $r = 0.061$, a choice of constants known to give rise to self-replicating RDSTs, or spots [40]. With this configuration, the spots are distinguished by local excess concentrations of B (relative to the equilibrium concentration of 0), collocated with local deficits in the concentration of A (relative to the equilibrium concentration of 1). We use the concentration of B as an order parameter for visual representations.

To couple this GSM to the chemical networks described in the previous section, we use reactions that are identical to those of the 0D case, with the Gray-Scott B component playing the role of the B component in the 0D version. The only change we applied to the chemical reactions are for the associative network \mathcal{N}_A , where the production of short-term memory M in the presence of S now occurs only in the presence of B. This is to reflect the fact that short-term memory can only be produced where spots are located:



1. Experimental relevance

In this numerical work, we used nondimensional units (for time, space, and concentration) to facilitate straightforward analysis and comparison with existing literature on numerical simulations of RDSs. However, we can easily translate these into meaningful experimental values. With a timescale of $dt = 1$ s, and a spatial scale of $dx = 10^{-2}$ cm, our simulation would represent a 2D patch of approximately 2.56×1.28 cm², with spots of approximately 1 mm evolving over timescales of ≈ 10 min. With those values, the diffusion coefficients would be of order 10^{-5} cm²/s. Such values are in line with experimental work that has been carried out on RDSs using, for instance, iodate-ferrocyanide-sulfite reactions in a thin layer of polyacrylamide gel [37].

2. Bolus delivery

In this spatial system, the delivery mode of the boluses of S and T is a key consideration and two different cases were explored. In the so-called *uniform* condition, boluses are delivered everywhere in a homogeneous fashion, as would occur for a 2D system sandwiched between porous plates, allowing the transverse diffusion of certain species (this is the way A is supplied and B removed in the classic GSM). Under this condition, all the spots come under equivalent attack from the toxin and are likewise homogeneously exposed to signals from the stimulus.

The second delivery mode is the *boundary diffusion* condition, in which boluses are delivered on the boundary of the domain, and permeate via diffusion. For this condition, we use higher bolus concentrations, along with larger diffusion

TABLE II. Diffusion constants, characteristic decay times, and reaction constants for the 2D networks. For stimulus and toxin diffusion coefficients, and stimulus bolus concentration, different values have been used for uniform and boundary diffusion conditions.

Parameter	Value	Description
C_S	1/10	Stimulus bolus concentration Uniform/Diffusion
D_S	1/20	Stimulus diffusion coefficient Uniform/Diffusion
D_T	1/20	Toxin diffusion coefficient Uniform/Diffusion
τ_S	100	Stimulus decay time
τ_T	100	Toxin decay time
D_N	1	Antidote diffusion coefficient
τ_N	400	Antidote decay time
k_{TB}	0.02	Reaction constant T+B
k_{NT}	1	Reaction constant N+T
k_{NB}	0.005	Reaction constant N+B
k_M	1	Short-term memory production reaction rate
k_L	0.025	Long-term memory production reaction rate
D_M	0.2	Short-term memory diffusion coefficient
τ_M	20	Short-term memory decay time
D_L	0.2	Long-term memory diffusion coefficient
τ_L	400	Long-term memory decay time

coefficients, to ensure that sufficient stimulus and toxin levels reach the center of the domain (see Table II). For both conditions, we used a delivery period of $\mathcal{T} = 100$ with an interval $\Delta\mathcal{T} = 50$ between stimulus and toxin. The concentration of A is initialised homogeneously to $\psi_A^{t=0} = 1$. To provide the initial seeds for the development of the spots, for each grid point $\psi_B^{t=0} = 1$ with a probability of 0.35, and $\psi_B^{t=0} = 0$ otherwise. To ensure the spots have time to fully develop before the toxin arrives, the boluses arrive only after 2000 time steps. Similar to the 0D case, the environment parameter $\epsilon(t)$, representing the magnitude of the toxin boluses relative to the stimulus boluses, rises linearly from 0 to 1 for 5000 time steps (from $t = 2000$ to $t = 7000$), then decreases linearly at the same rate (from $t = 7000$ to $t = 12000$) and remains at 0 thereafter.

The parameters of the reaction network were manually adjusted from the values found for the 0D case, and the three networks were compared at identical reaction rates (0.15 in the uniform case and 0.1 in the boundary diffusion case). Simulations were performed using forward Euler integration (we used $dt = 0.1$ for the uniform case, and a smaller value $dt = 0.01$ for the diffusion case, because of the larger diffusion coefficient). To speed up the core calculations, the simulation was implemented on a GPU using HLSL compute shaders, and Unity3D software for interface and rendering.

To quantify the evolution of the spot system while being perturbed by T, we computed the spatial average of ψ_B as a function of time. The number of spots was also counted using a thresholding and flood-fill procedure applied to the order parameter $\psi_B - \psi_A$ (threshold = -0.3).

B. Results

For both modes of bolus delivery (uniform and boundary diffusion), we simulated the three networks \mathcal{N}_D , \mathcal{N}_P , and \mathcal{N}_A in addition to a control case with an empty network (no antidote production). Results for the uniform and boundary

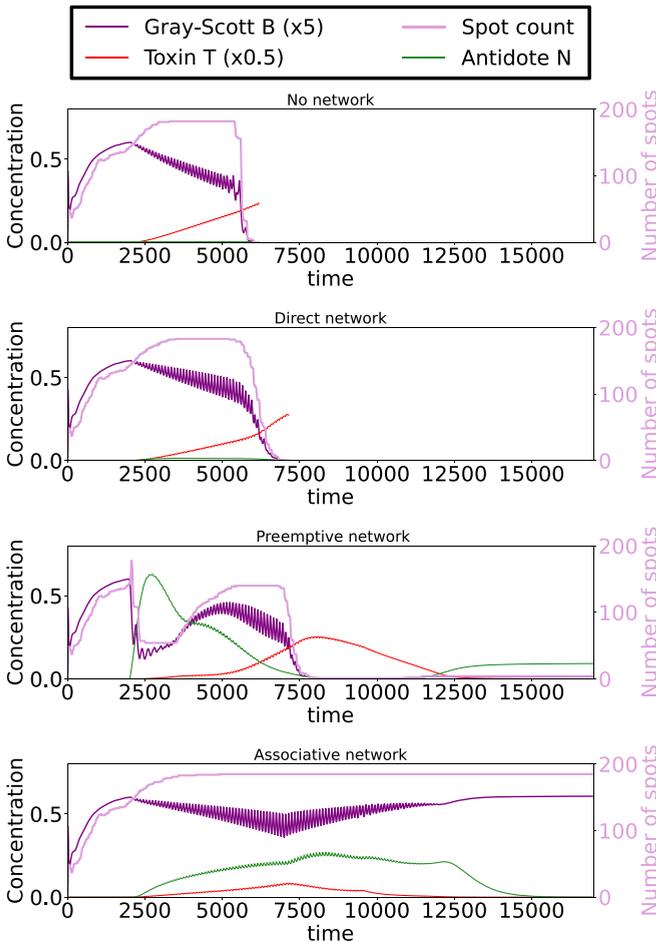


FIG. 7. Dynamics of uniform delivery system. Time series of the spatial average concentration of Gray Scott component B , Toxin T and Antidote N , along with the number of RDSTs, in the case of uniform bolus delivery: (1) no network, (2) direct network, (3) preemptive network, (4) associative learning network (time series of toxin and antidote concentration have been filtered to remove high frequency oscillations). See Table I for a summary of chemical equations used in each network.

delivery modes are displayed in Figs. 7 and 8, respectively. Full video of all eight simulations can be found in Ref. [162]. For both delivery mechanisms, we can draw similar conclusions, which are analogous to what was observed in the 0D case. In the control case (without any antidote production), the spots get quickly destroyed by the toxin. Oscillations of the average concentration of B can be observed, as toxin is periodically delivered and progressively damages the system, which only partially recovers in between boluses. This is analogous to the decreasing oscillations of B that were observed in the 0D case (see Fig. 2).

For the direct network \mathcal{N}_D , antidote is released when toxin boluses occur, but at an insufficient rate to prevent the complete destruction of the spots. As occurred in the 0D case, the system barely survives longer than when no antidote is produced.

For the preemptive network \mathcal{N}_P , antidote gets released when stimulus occurs, which constitutes an effective defense against the toxin. However, since it is only correlated with the

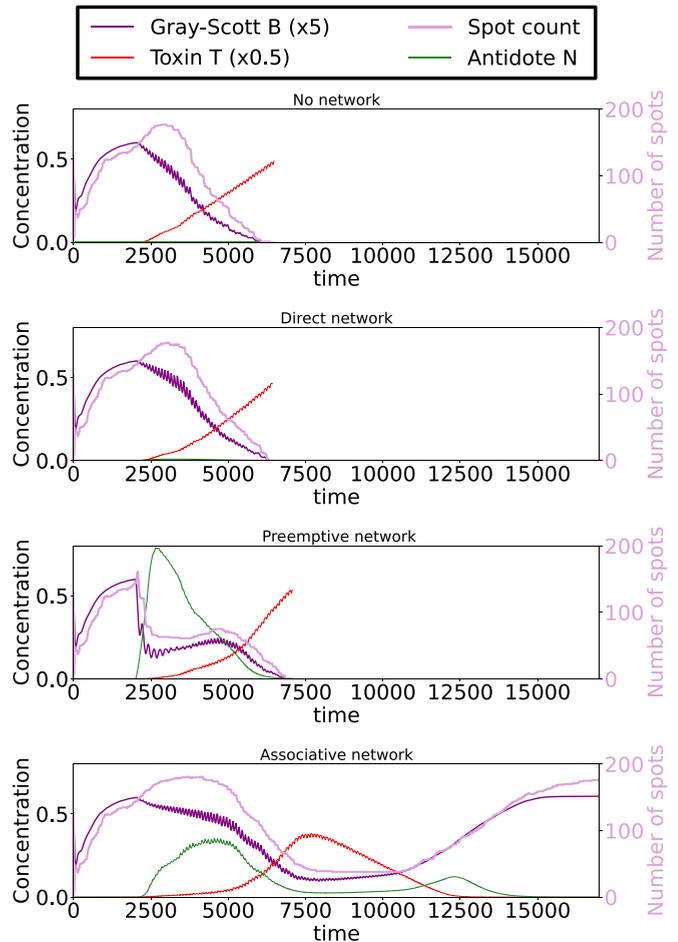


FIG. 8. Dynamics of boundary delivery system. Time series of the spatial average concentration of Gray Scott component B , Toxin T and Antidote N , along with the number of RDSTs, in the case of boundary diffusion bolus delivery: (1) no network, (2) direct network, (3) preemptive network, (4) associative learning network (time series of toxin and antidote concentration have been filtered to remove high frequency oscillations). See Table I for a summary of chemical equations used in each network.

stimulus level, the response is initially too strong when the toxin level is low, causing initial damage to the spots, and making them eventually disappear when the toxin level increases and causes additional damage, which is insufficiently mitigated by antidote production.

Finally, for the associative learning network \mathcal{N}_A , production of antidote occurs only when and where it is necessary. In this condition, from the count of spots and the average concentration of B (bottom left of Figs. 7 and 8), we can see that the spots get initially damaged by the toxin, until the progressive buildup of long-term memory triggers the appropriate response and prevents further destruction. The spots then replicate autocatalytically and reinhabit the space they previously lost (see Figs. 9 and 10 for snapshots of their spatial evolution). Overall, the features observed in the 0D-case are recapitulated in this spatial version. For the direct network, since antidote production is catalyzed by the presence of toxin, the response will naturally be tuned to the magnitude of the attack. However, the lack of an anticipation mechanism

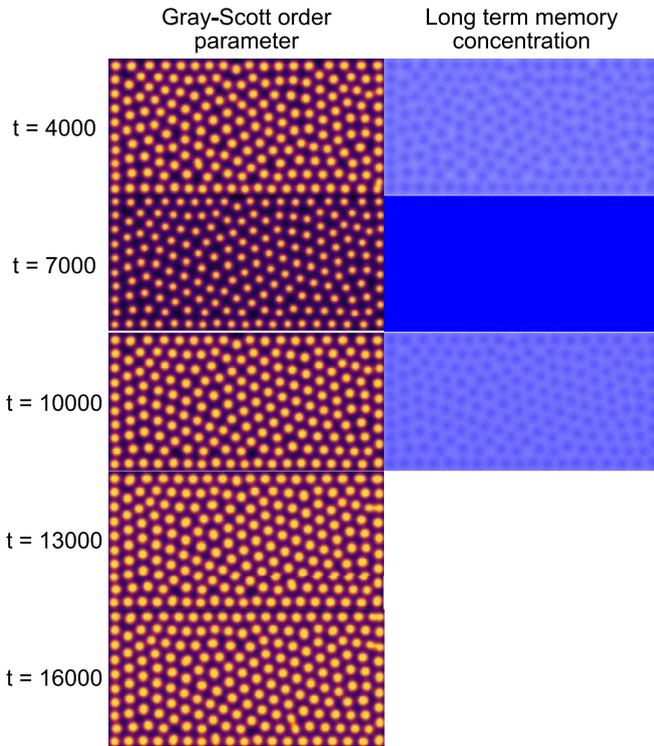


FIG. 9. Spatial behavior of uniform delivery system. Snapshots of the Gray-Scott B component and long-term memory L concentration at regular time intervals in the uniform delivery case, for the associative learning network \mathcal{N}_A . This sequence corresponds to the bottom panel of Fig. 7.

means that such a network would work only with a much higher reaction constant involving a very high instantaneous production rate of antidote. The preemptive network is able to produce antidote before toxin occurs but at a lower rate. This response is only proportionate to the level of stimulus, not the toxin. Hence, it lacks the ability to tune itself to the actual magnitude of the toxin concentration. A sensitivity analysis revealed that no value of the reaction rate k_p can lead to a survival of the system. A higher reaction constant is actually detrimental in the initial stages, when stimulus is high and toxin low, while a lower reaction constant leads to better survival at the beginning but increased damage when the toxin level becomes higher.

In contrast, we can see that the associative learning network combines both the anticipation mechanism of the preemptive network, and the self tuning ability of the direct network. A sensitivity analysis revealed that the network performs equally well for almost any value of the reaction rate k_D . As long as it is above a minimum of approximately 0.08 (for the uniform case), survival of the system occurs.

To illustrate the self-tuning of the associative network, and how it impacts its spatial differentiation abilities, it is instructive to compare it with the preemptive network. Figure 11 shows a concentration map after 10 toxin deliveries. We can see that for the associative network, antidote production is located on the boundary, as dictated by long-term memory concentration. However, for the preemptive network, antidote is produced by all the spots at similar rates, irrespec-

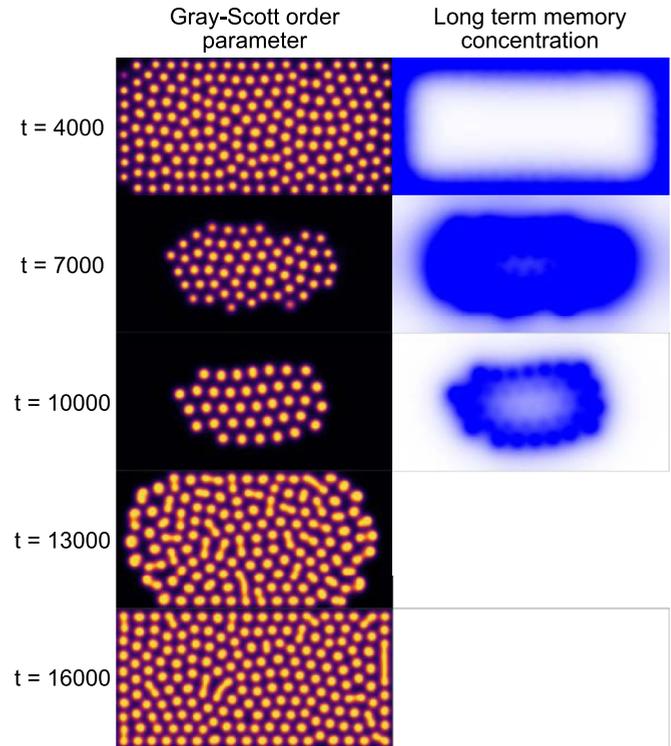


FIG. 10. Spatial behavior of boundary delivery system. Snapshots of the Gray-Scott B component and long-term memory L concentration at regular time intervals in the boundary diffusion delivery case, for the associative learning network \mathcal{N}_A . This sequence corresponds to the bottom panel of Fig. 8.

tive of their location. On the boundary, the antidote level is too low to resist high toxin levels. In the center, that same amount of antidote causes spots damage, since there is no toxin to degrade, and excess antidote is detrimental.

V. DISCUSSION

The results of the previous sections demonstrate a rudimentary associative learning ability of simple catalytic reaction networks. This was achieved using a pair of memory compounds with different characteristic timescales allowing external correlations to be encoded. In addition, the synthesis of long-term memory effectively proceeded through a logical AND operation on the presence of short-term memory and the presence of toxin. Hence, simple catalysis, an appropriate separation of timescales, and associative synthesis reactions appear to be the necessary and sufficient conditions for chemical associative learning [104,159].

Due to the associative feedback mechanism that catalyzes the production of long-term memory L , the system exhibits an automatic adaptation to the strength of the association between S and T (controlled by the environmental parameter ϵ). The level of L essentially becomes the learned estimation of the strength of this association. Hence, L is a discrete, low-pass filtered version of the time evolution of $\epsilon(t)$. The cut-off frequency of this filter is governed by the characteristic timescales of the network.

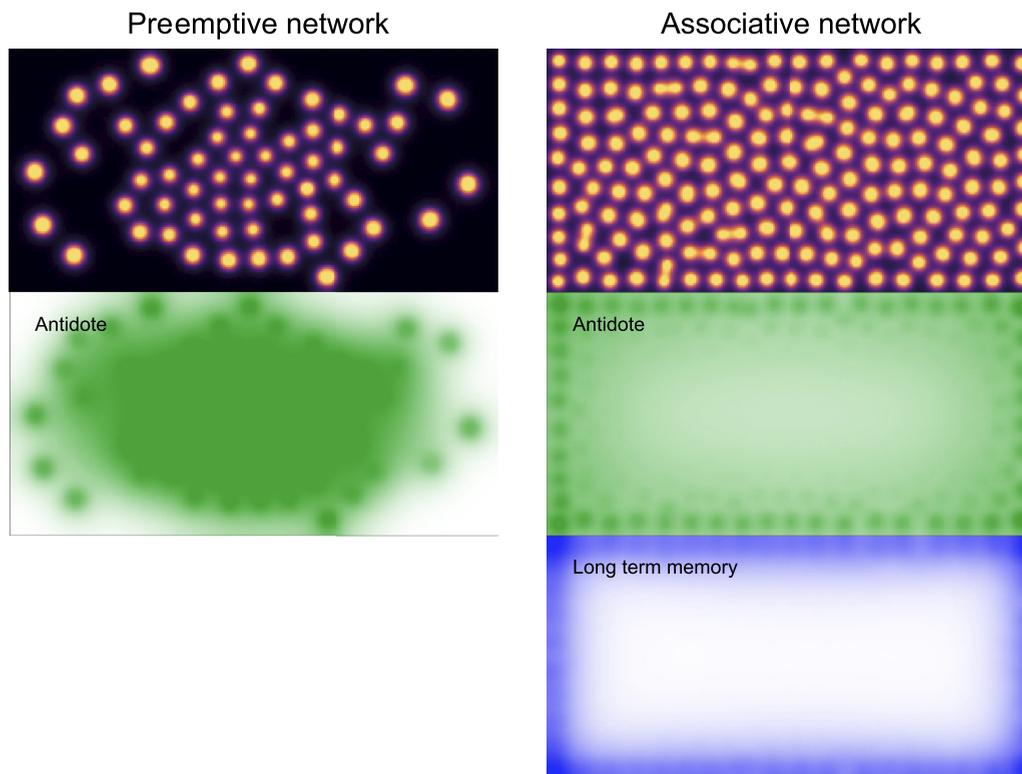


FIG. 11. Comparison between preemptive and associative concentration fields. Concentration map for preemptive and associative network in the boundary diffusion case, at the same time stamp ($t = 3010$, after 10 toxin deliveries.) For the associative network (right), antidote production is located at the boundary, as dictated by long-term memory concentration. For the preemptive network (left), antidote is produced by all spots in comparable amounts.

Tuning of our system parameters is based only on the temporal dynamics of the environment (period of delivery, spacing between stimulus and toxin, decay rate of the toxin, characteristic time of environmental changes). When all these factors are combined, only one parameter remains, and governs the tradeoff between residual toxin level and the amount of antidote produced. This parameter can be tuned based on the relative effect of toxin damage and antidote cost.

Note that in addition to lacking a learning ability, the preemptive network requires tuning to the strength of the toxin, whereas the associative network is self-tuning. The associative network could be less resilient if the toxin delivery events were short and fast, but there would be no favorability for very slow decay of long-term memory, so this only holds under extreme parameter values. In this work, the focus is situations beyond simple responses, where there is a tangible value in anticipation and learning. For scenarios where learning has no selective advantage, we would of course not expect any learning systems to emerge, so those scenarios are not relevant to the present work.

Beyond illustrating chemical associative learning, we presented simulations of emergent dissipative structures that can exploit these learning abilities to defend themselves against toxic compounds. These RDSTs exhibit a diverse repertoire of behaviors including replication (autocatalysis) [40], interstructure competition [36], homeostasis and symbiosis [34,35], among others [39,53,107,111,112,120,122]. With the results presented herein, we can now add associative learning

to the remarkable list of lifelike properties exhibited by RDSTs. In fact, they now satisfy the elementary four conditions to be classed as “lyfe” [77].

The addition of a simple learning network allows these emergent patterns to exploit environmental correlations to encode information (a single variable in this case), and use that information to adapt their behavior. Figures 9 and 10 illustrate how toxin delivery initially degrades the spot population until long-term memory builds up. This promotes the synthesis of antidote compound, and the spots then replicate into regions from which they had previously been eliminated.

In addition, the system shows a degree of functional differentiation, in which the spots that are exposed to toxin and stimulus build up a local memory and produce antidote, whereas isolated spots in the core of the group are largely shielded by the toxin (in the boundary delivery case, which is spatially heterogeneous in terms of the delivery of toxin and stimulus). This illustrates that specialization in a colony of dissipative structures naturally follows from differences in environmental selective forces, and does not require genetically-encoded, cellular structures or complex molecular machinery. A similar spatial division of function was observed in thermal homeostasis models of RDSTs [35]. In the context of prebiotic structure formation, one could even speculate that environmental spatial heterogeneity may be the driving force for the emergence of individuated patterns. Perhaps the differing physical selective forces placed upon such heterogeneous systems naturally promotes the emergence of boundary structures.

A natural question that arises from our results is the degree to which the emergent behaviors are genuinely emergent. One might contend that this system is simply a hand-designed reaction network that necessarily achieves the desired goal because it was engineered to do so. Thus how does this have any relevance to the natural world? The relevance of our results lies in the demonstration that such a simple chemical learning mechanism is possible and the conditions under which it is dynamically favored. What we have presented is a locus from which the feasibility of emergent learning in OoL contexts can be explored. Having demonstrated the core requirements for autonomous chemical learning, the next question to answer is how such behavior would still be favored in the presence of larger networks populated by deleterious side reactions that are neutral or detrimental to learning? This would constitute a rigorous test of the emergence of the demonstrated learning behaviors.

Specifically, we intend to embed our system in progressively larger networks containing random reactions. The null hypothesis of such an investigation would be that adding superfluous reactions to learning networks disrupts and obscures the learning behavior, and the system simply loses its learning abilities. Rejection of such a hypothesis would require the observation of conditions in which learning is dynamically favored by a set of external conditions, i.e., a positive feedback between learning and the resilience or persistence of the system performing the learning. In this case there would be a subtle interplay between thermodynamic, kinetic and higher level effects that is difficult to predict without numerical or experimental modeling. The second law of thermodynamics is of limited utility because the system is not isolated and short-term or stochastic effects can be amplified nonlinearly. However, the system's dynamics on average would be exergonic. Hence, our next phase of modeling will explore the selection of learning behavior from a larger, naturally disordered space of dynamical processes.

We presented these results in the context of a protometabolic system that is vulnerable to degradation by a toxic compound. This compound's arrival is time varying but predictable via the occurrence of its precursor compound in associated environments. While appearing somewhat esoteric, this scenario could be generalized to other situations, such as predicting the arrival of nutrient compounds, or any exploitable association of time-varying events. If a learning system existed in an environment with no learnable features, then clearly a learning ability would be futile, offer no selective advantage, and would not persist (instantaneous feedback responses with no memory would be the optimal strategies in that case). However, life is surrounded by learnable features, and despite the stochastic and challenging nature of the Hadean and Archaean eons, there would still have been periodic, correlated, or predictable features of the environment that a nascent prebiotic system could potentially exploit (at the very least, diurnal or tidal cycles that would correlate with external fluxes of different compounds).

The primary feature separating life from the abiotic world is its unique information-processing abilities [134–141,143–150], whereas nonliving entities are at best only weakly predictive, teleological, nor do they exhibit perception-action cycles or measurement-feedback protocols

[163–165]. When considering the emergence of the living state, it seems clear that information-processing and learning are a fundamental feature that must arise at some point. Given its importance, we suggest that that point was relatively early in the abiogenesis story, and we hope that the results presented here may provide a fruitful path to understanding that emergence in a general context.

In the OoL field, tremendous progress has been achieved in the areas of droplet systems [15,16,29–33] and coacervates [43,59,60,166–170]. In both cases, nonlipid molecules spontaneously produce individuated structures or boundaries. Such structures readily exhibit complex behavior such as self-replication, chemotaxis, molecular concentration, or the enhancement of oligimerization reactions. We suggest that the incorporation of simple learning circuits as presented in this work could add a novel and essential feature to these protocellular systems. In addition, it is likely that our learning circuits could be implemented using peptide or RNA-based systems. This could add a whole new dimension to contemporary origins experiments.

VI. CONCLUSIONS

This work presents a simple catalytic reaction network capable of encoding correlations in environmental variables. In the context of emergent, prebiotic systems, such an associative learning ability could provide a defense mechanism against detrimental, but predictable perturbations imposed by an unstable environment. We demonstrated how associative chemical learning can be implemented using long- and short-term memory compounds. This allows the system to sense its environment over different timescales and encode the presence of correlated features. The synthesis of long-term memory effectively implements an AND operation on the short-term memory and the environmental variable being sensed. Such a mechanism allows the long-term memory concentration to serve as a low-pass filter on the environmental variable (the degree to which a toxin is correlated with a precursor compound). Hence, the learning ability is continuous, as opposed to discrete.

We showed that such a learning capacity can enhance the resilience of emergent chemical patterns (RDSTs). When endowed with learning networks, these structures are capable of autocatalysis, homeostasis and rudimentary learning. Implementations of this chemical learning system in OoL models such as droplets, coacervates or mineral surfaces could introduce a novel level of adaptive functionality that is an essential step in the journey from geophysics and geochemistry to biophysics, biochemistry, and information processing.

Future work will carry out rigorous tests on the persistence of learning in the presence of deleterious side reactions. Establishing these necessary and sufficient conditions for emergent learning would have significant implications for research in Artificial life, OoL, systems chemistry, complexity science, and synthetic biology.

ACKNOWLEDGMENTS

We gratefully acknowledge Artemy Kolchinsky of the University of Tokyo and Sante Fe Institute for his insightful and constructive review of a prepublication version of this

manuscript. This work was supported by the Caltech Division of Geological and Planetary Sciences Discovery Fund. We thank the Caltech GPS “Astrobiothermoinfoevo” group for the numerous inspiring discussions that helped catalyze this work.

Similarly, we thank the various members of the Earth-Life Science Institute and EON program members for creating such an inspiring and creative environment for intellectual exploration.

-
- [1] D. Kondepudi and I. Prigogine, *Modern Thermodynamics: From Heat Engines to Dissipative Structures* (Wiley, New York, NY, 2014).
- [2] I. Prigogine, Time, structure, and fluctuations, *Science* **201**, 777 (1978).
- [3] P. W. Anderson *et al.*, More is different, *Science* **177**, 393 (1972).
- [4] J. Goldstein, Emergence as a construct: History and issues, *Emergence* **1**, 49 (1999).
- [5] E. P. Hoel, L. Albantakis, and G. Tononi, Quantifying causal emergence shows that macro can beat micro, *Proc. Natl. Acad. Sci. USA* **110**, 19790 (2013).
- [6] J. Holland, *Emergence: From Chaos to Order*, Popular Science/Oxford University Press (Oxford University Press, Oxford, UK, 2000).
- [7] S. Rasmussen, N. A. Baas, B. Mayer, M. Nilsson, and M. W. Olesen, Ansatz for dynamical hierarchies, *Artific. Life* **7**, 329 (2001).
- [8] G. Ahlers, S. Grossmann, and D. Lohse, Heat transfer and large scale dynamics in turbulent Rayleigh-Bénard convection, *Rev. Mod. Phys.* **81**, 503 (2009).
- [9] S. Bartlett and S. Bullock, Natural convection of a two-dimensional Boussinesq fluid does not maximize entropy production, *Phys. Rev. E* **90**, 023014 (2014).
- [10] S. Bartlett and N. Virgo, Maximum entropy production is not a steady state attractor for 2D fluid convection, *Entropy* **18**, 431 (2016).
- [11] P. Manneville, Rayleigh-Bénard convection: Thirty years of experimental, theoretical, and modeling work, in *Dynamics of Spatio-Temporal Cellular Structures: Henri Bénard Centenary Review* (Springer, New York, 2006), pp. 41–65.
- [12] W. Pesch, Complex spatiotemporal convection patterns, *Chaos* **6**, 348 (1996).
- [13] K. A. Emanuel, The maximum intensity of hurricanes, *J. Atmos. Sci.* **45**, 1143 (1988).
- [14] P. A. Bachmann, P. L. Luisi, and J. Lang, Autocatalytic self-replicating micelles as models for prebiotic structures, *Nature (London)* **357**, 57 (1992).
- [15] M. M. Hanczyc, Metabolism and motility in prebiotic structures, *Philos. Trans. R. Soc. B* **366**, 2885 (2011).
- [16] M. M. Hanczyc, T. Toyota, T. Ikegami, N. Packard, and T. Sugawara, Fatty acid chemistry at the oil-water interface: Self-propelled oil droplets, *J. Am. Chem. Soc.* **129**, 9386 (2007).
- [17] B. Mayer, G. Köhler, and S. Rasmussen, Simulation and dynamics of entropy-driven, molecular self-assembly processes, *Phys. Rev. E* **55**, 4489 (1997).
- [18] D. Segré, D. Ben-Eli, D. W. Deamer, and D. Lancet, The lipid world, *Origins Life Evol. Biosphere* **31**, 119 (2001).
- [19] T. F. Zhu and J. W. Szostak, Coupled growth and division of model protocell membranes, *J. Am. Chem. Soc.* **131**, 5705 (2009).
- [20] T. Nozakura and S. Ikeuchi, Formation of dissipative structures in galaxies, *Astrophys. J.* **279**, 40 (1984).
- [21] A. Curir, Rotating black holes as dissipative spin-thermodynamical systems, *Gen. Relativ. Gravit.* **13**, 417 (1981).
- [22] D. Kondepudi, B. Kay, and J. Dixon, Dissipative structures, machines, and organisms: A perspective, *Chaos* **27**, 104607 (2017).
- [23] C. H. Lineweaver and C. A. Egan, Life, gravity, and the second law of thermodynamics, *Phys. Life Rev.* **5**, 225 (2008).
- [24] N. D. Virgo, Thermodynamics and the structure of living systems, Ph.D. thesis, University of Sussex Brighton (2011).
- [25] B. J. Cafferty, A. S. Wong, S. N. Semenov, L. Belding, S. Gmür, W. T. Huck, and G. M. Whitesides, Robustness, entrainment, and hybridization in dissipative molecular networks, and the origin of life, *J. Am. Chem. Soc.* **141**, 8289 (2019).
- [26] M. Egbert, Y. Kolezhitskiy, and N. Virgo, Steering the growth of adaptive self-preserving dissipative structures, in *Artificial Life Conference Proceedings* (MIT Press, Cambridge, MA, 2019), pp. 255–262.
- [27] T. Froese, N. Virgo, and T. Ikegami, Motility at the origin of life: Its characterization and a model, *Artific. Life* **20**, 55 (2014).
- [28] D. Lancet, R. Zidovetzki, and O. Markovitch, Systems proto-biology: Origin of life in lipid catalytic networks, *J. R. Soc. Interface* **15**, 20180159 (2018).
- [29] J. Čejková, M. Novak, F. Štěpánek, and M. M. Hanczyc, Dynamics of chemotactic droplets in salt concentration gradients, *Langmuir* **30**, 11937 (2014).
- [30] J. Čejková, S. Holler, T. Q. Nguyenová, C. Kerrigan, F. Štěpánek, and M. M. Hanczyc, Chemotaxis and chemokinesis of living and nonliving objects, in *Advances in Unconventional Computing* (Springer, Berlin, 2017), pp. 245–260.
- [31] J. Čejková, T. Banno, M. M. Hanczyc, and F. Štěpánek, Droplets as liquid robots, *Artific. Life* **23**, 528 (2017).
- [32] M. M. Hanczyc, S. M. Fujikawa, and J. W. Szostak, Experimental models of primitive cellular compartments: Encapsulation, growth, and division, *Science* **302**, 618 (2003).
- [33] S. Holler and M. M. Hanczyc, Droplet based synthetic biology: Chemotaxis and interface with biology, in *Artificial Life Conference Proceedings* (MIT Press, Cambridge, MA, 2019), pp. 650–651.
- [34] S. Bartlett, Why is life? An assessment of the thermodynamic properties of dissipative, pattern-forming systems, Ph.D. thesis, University of Southampton (2014).
- [35] S. Bartlett and S. Bullock, A precarious existence: Thermal homeostasis of simple dissipative structures, in *Proceedings of the 15th International Conference on the Synthesis and Simulation of Living Systems* (MIT Press, Cambridge, MA, 2016), pp. 608–615.
- [36] S. Bartlett and S. Bullock, Emergence of competition between different dissipative structures for the same free energy source,

- in *Proceedings of the Artificial Life Conference Proceedings 13* (MIT Press, Cambridge, MA, 2015), pp. 415–422.
- [37] K.-J. Lee, W. D. McCormick, J. E. Pearson, and H. L. Swinney, Experimental observation of self-replicating spots in a reaction–diffusion system, *Nature (London)* **369**, 215 (1994).
- [38] K. J. Lee and H. L. Swinney, Lamellar structures and self-replicating spots in a reaction-diffusion system, *Phys. Rev. E* **51**, 1899 (1995).
- [39] F. Lesmes, D. Hochberg, F. Morán, and J. Pérez-Mercader, Noise-Controlled Self-Replicating Patterns, *Phys. Rev. Lett.* **91**, 238301 (2003).
- [40] J. E. Pearson, Complex patterns in a simple system, *Science* **261**, 189 (1993).
- [41] I. A. Chen, R. W. Roberts, and J. W. Szostak, The emergence of competition between model protocells, *Science* **305**, 1474 (2004).
- [42] M. M. Hanczyc and J. W. Szostak, Replicating vesicles as models of primitive cell growth and division, *Curr. Opin. Chem. Biol.* **8**, 660 (2004).
- [43] T. Z. Jia, K. Chandru, Y. Hongo, R. Afrin, T. Usui, K. Myojo, and H. J. Cleaves, Membraneless polyester microdroplets as primordial compartments at the origins of life, *Proc. Natl. Acad. Sci. USA* **116**, 15830 (2019).
- [44] S. F. Jordan, H. Ramm, I. N. Zheludev, A. M. Hartley, A. Maréchal, and N. Lane, Promotion of protocell self-assembly from mixed amphiphiles at the origin of life, *Nat. Ecol. Evol.* **3**, 1705 (2019).
- [45] G. F. Joyce and J. W. Szostak, Protocells and RNA self-replication, *Cold Spring Harbor Perspect. Biol.* **10**, a034801 (2018).
- [46] S. Mann, Systems of creation: The emergence of life from nonliving matter, *Acc. Chem. Res.* **45**, 2131 (2012).
- [47] N. Ono, Computational studies on conditions of the emergence of autopoietic protocells, *Biosystems* **81**, 223 (2005).
- [48] S. Rasmussen, A. Constantinescu, and C. Svaneborg, Generating minimal living systems from non-living materials and increasing their evolutionary abilities, *Philos. Trans. R. Soc. B* **371**, 20150440 (2016).
- [49] B. Shirt-Ediss, R. V. Solé, and K. Ruiz-Mirazo, Emergent chemical behavior in variable-volume protocells, *Life* **5**, 181 (2015).
- [50] R. V. Solé, A. Munteanu, C. Rodriguez-Caso, and J. Macía, Synthetic protocell biology: From reproduction to computation, *Philos. Trans. R. Soc. B* **362**, 1727 (2007).
- [51] K. Adamala and J. W. Szostak, Competition between model protocells driven by an encapsulated catalyst, *Nat. Chem.* **5**, 495 (2013).
- [52] A. E. Engelhart, K. P. Adamala, and J. W. Szostak, A simple physical mechanism enables homeostasis in primitive cells, *Nat. Chem.* **8**, 448 (2016).
- [53] T. Froese, T. Ikegami, and N. Virgo, The behavior-based hypercycle: From parasitic reaction to symbiotic behavior, *Artific. Life* **13**, 457 (2012).
- [54] H. R. Maturana and F. J. Varela, *Autopoiesis and Cognition: The Realization of the Living* (Springer Science & Business Media, 2012), Vol. 42.
- [55] T. Gánti, Organization of chemical reactions into dividing and metabolizing units: The chemotons, *Biosystems* **7**, 15 (1975).
- [56] S. Rasmussen, M. A. Bedau, J. S. McCaskill, and N. H. Packard, A roadmap to protocells, in *Protocells: Bridging Nonliving and Living Matter* (MIT Press, Cambridge, MA, 2008), pp. 71–100.
- [57] A. Joesaar, S. Yang, B. Bögers, A. van der Linden, P. Pieters, B. P. Kumar, N. Dalchau, A. Phillips, S. Mann, and T. F. de Greef, DNA-based communication in populations of synthetic protocells, *Nat. Nanotechnol.* **14**, 369 (2019).
- [58] K. Kurihara, Y. Okura, M. Matsuo, T. Toyota, K. Suzuki, and T. Sugawara, A recursive vesicle-based model protocell with a primitive model cell cycle, *Nat. Commun.* **6**, 8352 (2015).
- [59] M. Li, X. Huang, T. D. Tang, and S. Mann, Synthetic cellularity based on non-lipid micro-compartments and protocell models, *Curr. Opin. Chem. Biol.* **22**, 1 (2014).
- [60] Y. Qiao, M. Li, R. Booth, and S. Mann, Predatory behavior in synthetic protocell communities, *Nat. Chem.* **9**, 110 (2017).
- [61] A. Aubrey, H. Cleaves, and J. L. Bada, The role of submarine hydrothermal systems in the synthesis of amino acids, *Orig. Life Evol. Biosph.* **39**, 91 (2009).
- [62] K. Chandru, A. Gilbert, C. Butch, M. Aono, and H. J. Cleaves, The abiotic chemistry of thiolated acetate derivatives and the origin of life, *Sci. Rep.* **6**, 29883 (2016).
- [63] R. Krishnamurthy, Giving rise to life: Transition from prebiotic chemistry to protobiology, *Acc. Chem. Res.* **50**, 455 (2017).
- [64] R. Krishnamurthy, Life’s biological chemistry: A destiny or destination starting from prebiotic chemistry? *Chem. Eur. J.* **24**, 16708 (2018).
- [65] S. B. Carroll, Chance and necessity: The evolution of morphological complexity and diversity, *Nature (London)* **409**, 1102 (2001).
- [66] S. B. Carroll, *A Series of Fortunate Events: Chance and the Making of the Planet, Life, and You* (Princeton University Press, Princeton, NJ, 2020).
- [67] S. J. Gould and E. S. Vrba, Exaptation—A missing term in the science of form, *Paleobiology* **8**, 4 (1982).
- [68] K. Chandru, I. Mamajanov, H. J. Cleaves, and T. Z. Jia, Polyesters as a model system for building primitive biologies from non-biological prebiotic chemistry, *Life* **10**, 6 (2020).
- [69] K. Chandru, T. Z. Jia, I. Mamajanov, N. Bapat, and H. J. Cleaves, Prebiotic oligomerization and self-assembly of structurally diverse xenobiological monomers, *Sci. Rep.* **10**, 17560 (2020).
- [70] C. Cleland and C. Chyba, Does life have a definition? in *Planets and Life: The Emerging Science of Astrobiology* (Cambridge University Press, Cambridge, UK, 2007), pp. 119–131.
- [71] A. Cornish-Bowden and M. L. Cárdenas, Contrasting theories of life: Historical context, current theories. In search of an ideal theory, *Biosystems* **188**, 104063 (2020).
- [72] C. Mariscal and W. F. Doolittle, Life and life only: A radical alternative to life definitionism, *Synthese* **197**, 2975 (2020).
- [73] L. J. Mix, Defending definitions of life, *Astrobiology* **15**, 15 (2015).
- [74] K. Ruiz-Mirazo, J. Peretó, and A. Moreno, A universal definition of life: Autonomy and open-ended evolution, *Orig. Life Evol. Biosph.* **34**, 323 (2004).
- [75] M. Vitas and A. Dobovišek, Toward a general definition of life, *Orig. Life Evol. Biosph.* **49**, 77 (2019).
- [76] G. Witzany, What is life? *Front. Astron. Space Sci.* **7**, 7 (2020).
- [77] S. Bartlett and M. L. Wong, Defining life in the universe: From three privileged functions to four pillars, *Life* **10**, 42 (2020).

- [78] R. Bar-Ziv, T. Tlusty, and A. Libchaber, Protein–DNA computation by stochastic assembly cascade, *Proc. Natl. Acad. Sci. USA* **99**, 11589 (2002).
- [79] Y. Benenson, Biomolecular computing systems: Principles, progress and potential, *Nat. Rev. Genet.* **13**, 455 (2012).
- [80] J. Elbaz, O. Lioubashevski, F. Wang, F. Remacle, R. D. Levine, and I. Willner, DNA computing circuits using libraries of DNAzyme subunits, *Nat. Nanotechnol.* **5**, 417 (2010).
- [81] L. Qian, E. Winfree, and J. Bruck, Neural network computation with DNA strand displacement cascades, *Nature (London)* **475**, 368 (2011).
- [82] E. Shapiro and Y. Benenson, Bringing DNA computers to life, *Sci. Am.* **294**, 44 (2006).
- [83] A. Adamatzky and B. De Lacy Costello, Experimental logical gates in a reaction-diffusion medium: The XOR gate and beyond, *Phys. Rev. E* **66**, 046112 (2002).
- [84] M. Dueñas-Díez and J. Pérez-Mercader, *In vitro* reconfigurability of native chemical automata, the inclusiveness of their hierarchy and their thermodynamics, *Sci. Rep.* **10**, 6814 (2020).
- [85] J. Gorecki, K. Gizynski, J. Guzowski, J. Gorecka, P. Garstecki, G. Gruenert, and P. Dittrich, Chemical computing with reaction–diffusion processes, *Philos. Trans. R. Soc. A* **373**, 20140219 (2015).
- [86] W. Hohmann, M. Kraus, and F. Schneider, Learning and recognition in excitable chemical reactor networks, *J. Phys. Chem. A* **102**, 3103 (1998).
- [87] J. Holley, I. Jahan, B. De Lacy Costello, L. Bull, and A. Adamatzky, Logical and arithmetic circuits in Belousov-Zhabotinsky encapsulated disks, *Phys. Rev. E* **84**, 056110 (2011).
- [88] J. M. Parrilla-Gutierrez, A. Sharma, S. Tsuda, G. J. T. Cooper, G. Aragon-Camarasa, K. Donkers, and L. Cronin, A programmable chemical computer with memory and pattern recognition, *Nat. Commun.* **11**, 1442 (2020).
- [89] O. Steinbock, P. Kettunen, and K. Showalter, Chemical wave logic gates, *J. Phys. Chem.* **100**, 18970 (1996).
- [90] M. Sun and X. Zhao, Combinational logic circuit based on BZ reaction, in *Advances in Unconventional Computing* (Springer, Berlin, 2017), pp. 105–139.
- [91] A. Wang, J. Gold, N. Tompkins, M. Heymann, K. Harrington, and S. Fraden, Configurable NOR gate arrays from Belousov-Zhabotinsky micro-droplets, *Eur. Phys. J.: Spec. Top.* **225**, 211 (2016).
- [92] W. Banzhaf, P. Dittrich, and H. Rauhe, Emergent computation by catalytic reactions, *Nanotechnology* **7**, 307 (1996).
- [93] W. Banzhaf and L. Yamamoto, *Artificial Chemistries* (MIT Press, Cambridge, MA, 2015).
- [94] D. Blount, P. Banda, C. Teuscher, and D. Stefanovic, Feedforward chemical neural network: An in silico chemical system that learns XOR, *Artific. Life* **23**, 295 (2017).
- [95] H.-L. Chen, D. Doty, and D. Soloveichik, Deterministic function computation with chemical reaction networks, *Nat. Comput.* **13**, 517 (2014).
- [96] M. Dueñas-Díez and J. Pérez-Mercader, How chemistry computes: Language recognition by non-biochemical chemical automata. from finite automata to turing machines, *iScience* **19**, 514 (2019).
- [97] J. Gasteiger and J. Zupan, Neural networks in chemistry, *Angew. Chem. Int. Ed. Engl.* **32**, 503 (1993).
- [98] A. Hjelmfelt, E. D. Weinberger, and J. Ross, Chemical implementation of neural networks and turing machines, *Proc. Natl. Acad. Sci. USA* **88**, 10983 (1991).
- [99] A. Hjelmfelt and J. Ross, Chemical implementation and thermodynamics of collective neural networks, *Proc. Natl. Acad. Sci. USA* **89**, 388 (1992).
- [100] A. Hjelmfelt, F. Schneider, and J. Ross, Pattern recognition in coupled chemical kinetic systems, *Science* **260**, 335 (1993).
- [101] A. Hjelmfelt and J. Ross, Implementation of logic functions and computations by chemical kinetics, *Physica D* **84**, 180 (1995).
- [102] M. O. Magnasco, Chemical Kinetics is Turing Universal, *Phys. Rev. Lett.* **78**, 1190 (1997).
- [103] F. Muzika and I. Schreiber, Control of turing patterns and their usage as sensors, memory arrays, and logic gates, *J. Chem. Phys.* **139**, 164108 (2013).
- [104] W. Poole, A. Ortiz-Munoz, A. Behera, N. S. Jones, T. E. Ouldrige, E. Winfree, and M. Gopalkrishnan, Chemical Boltzmann machines, in *International Conference on DNA-Based Computers* (Springer, Cham, 2017), pp. 210–231.
- [105] D. Soloveichik, M. Cook, E. Winfree, and J. Bruck, Computation with finite stochastic chemical reaction networks, *Nat. Comput.* **7**, 615 (2008).
- [106] K. Dale and P. Husbands, The evolution of reaction-diffusion controllers for minimally cognitive agents, *Artific. Life* **16**, 1 (2010).
- [107] A. Awazu and K. Kaneko, Relaxation to Equilibrium Can be Hindered by Transient Dissipative Structures, *Phys. Rev. Lett.* **92**, 258302 (2004).
- [108] M. Budroni and A. De Wit, Dissipative structures: From reaction-diffusion to chemo-hydrodynamic patterns, *Chaos* **27**, 104617 (2017).
- [109] I. R. Epstein, Predicting complex biology with simple chemistry, *Proc. Natl. Acad. Sci. USA* **103**, 15727 (2006).
- [110] I. R. Epstein and B. Xu, Reaction–diffusion processes at the nano- and microscales, *Nat. Nanotechnol.* **11**, 312 (2016).
- [111] J.-S. Gagnon, D. Hochberg, and J. Pérez-Mercader, Small-scale properties of a stochastic cubic-autocatalytic reaction-diffusion model, *Phys. Rev. E* **92**, 042114 (2015).
- [112] J.-S. Gagnon, D. Hochberg, and J. Pérez-Mercader, Selection and control of pathways by using externally adjustable noise on a stochastic cubic autocatalytic chemical system, *Phys. Rev. E* **98**, 062216 (2018).
- [113] P. Gray and S. Scott, Sustained oscillations and other exotic patterns of behavior in isothermal reactions, *J. Phys. Chem.* **89**, 22 (1985).
- [114] P. Gray, S. K. Scott, and J. H. Merkin, The Brusselator model of oscillatory reactions: Relationships between two-variable and four-variable models with rigorous application of mass conservation and detailed balance, *J. Chem. Soc. Faraday Trans. 1* **84**, 993 (1988).
- [115] J. Halatek and E. Frey, Rethinking pattern formation in reaction–diffusion systems, *Nat. Phys.* **14**, 507 (2018).
- [116] M. Kuznetsov, A. Kolobov, and A. Polezhaev, Pattern formation in a reaction-diffusion system of Fitzhugh-Nagumo type before the onset of subcritical turing bifurcation, *Phys. Rev. E* **95**, 052208 (2017).
- [117] Y. Kyrychko, K. Blyuss, S. Hogan, and E. Schöll, Control of spatiotemporal patterns in the Gray–Scott model, *Chaos* **19**, 043126 (2009).

- [118] K. J. Lee, W. McCormick, Q. Ouyang, and H. L. Swinney, Pattern formation by interacting chemical fronts, *Science* **261**, 192 (1993).
- [119] H. Mahara, K. Suzuki, R. A. Jahan, and T. Yamaguchi, Co-existing stable patterns in a reaction-diffusion system with reversible Gray-Scott dynamics, *Phys. Rev. E* **78**, 066210 (2008).
- [120] Y. Nishiura and D. Ueyama, Spatio-temporal chaos for the Gray-Scott model, *Physica D* **150**, 137 (2001).
- [121] H.-G. Purwins, H. Bödeker, and A. Liehr, Dissipative solitons in reaction-diffusion systems, in *Dissipative Solitons* (Springer, Berlin, 2005), pp. 267–308.
- [122] M. Stich, G. Ghoshal, and J. Pérez-Mercader, Parametric pattern selection in a reaction-diffusion model, *PLoS ONE* **8**, e77337 (2013).
- [123] A. M. Turing, The chemical basis of morphogenesis, *Philos. Trans. R. Soc. Lond. B* **237**, 37 (1952).
- [124] E. E. Sel'kov, Self-oscillations in glycolysis I. A simple kinetic model, *Eur. J. Biochem.* **4**, 79 (1968).
- [125] V. Horváth, I. R. Epstein, and K. Kustin, Mechanism of the ferrocyanide-iodate-sulfite oscillatory chemical reaction, *J. Phys. Chem. A* **120**, 1951 (2016).
- [126] I. Szalai and P. De Kepper, Patterns of the ferrocyanide-iodate-sulfite reaction revisited: The role of immobilized carboxylic functions, *J. Phys. Chem. A* **112**, 783 (2008).
- [127] I. Szalai and P. De Kepper, Pattern formation in the ferrocyanide-iodate-sulfite reaction: The control of space scale separation, *Chaos* **18**, 026105 (2008).
- [128] M. Bertram, C. Beta, M. Pollmann, A. S. Mikhailov, H. H. Rotermund, and G. Ertl, Pattern formation on the edge of chaos: Experiments with co oxidation on a pt (110) surface under global delayed feedback, *Phys. Rev. E* **67**, 036208 (2003).
- [129] C. Beta, M. Bertram, A. S. Mikhailov, H. H. Rotermund, and G. Ertl, Controlling turbulence in a surface chemical reaction by time-delay autosynchronization, *Phys. Rev. E* **67**, 046224 (2003).
- [130] G. Ertl, Oscillatory kinetics and spatio-temporal self-organization in reactions at solid surfaces, *Science* **254**, 1750 (1991).
- [131] G. Ertl, Reactions at surfaces: From atoms to complexity (nobel lecture), *Angew. Chem. Int. Ed.* **47**, 3524 (2008).
- [132] Y.-J. Li, J. Osolonovitch, N. Mazouz, F. Plenge, K. Krischer, and G. Ertl, Turing-type patterns on electrode surfaces, *Science* **291**, 2395 (2001).
- [133] C. Sachs, M. Hildebrand, S. Völkening, J. Wintterlin, and G. Ertl, Spatiotemporal self-organization in a surface reaction: From the atomic to the mesoscopic scale, *Science* **293**, 1635 (2001).
- [134] F. Baluška and M. Levin, On having no head: Cognition throughout biological systems, *Front. Psychol.* **7**, 902 (2016).
- [135] E. Ben-Jacob, Learning from bacteria about natural information processing, *Ann. N.Y. Acad. Sci.* **1178**, 78 (2009).
- [136] E. Ben-Jacob, M. Lu, D. Schultz, and J. Onuchic, The physics of bacterial decision making, *Front. Cell. Infect. Microbiol.* **4**, 154 (2014).
- [137] J. Delgado and R. V. Solé, Collective-induced computation, *Phys. Rev. E* **55**, 2338 (1997).
- [138] Z. Erez, I. Steinberger-Levy, M. Shamir, S. Doron, A. Stokar-Avihail, Y. Peleg, S. Melamed, A. Leavitt, A. Savidor, S. Albeck *et al.*, Communication between viruses guides lysis-lysogeny decisions, *Nature (London)* **541**, 488 (2017).
- [139] K. D. Farnsworth, J. Nelson, and C. Gershenson, Living is information processing: From molecules to global systems, *Acta Biotheor.* **61**, 203 (2013).
- [140] S. J. Gershman, P. E. Balbi, C. R. Gallistel, and J. Gunawardena, Reconsidering the evidence for learning in single cells, *Elife* **10**, e61907 (2021).
- [141] S. Ginsburg and E. Jablonka, The evolution of associative learning: A factor in the cambrian explosion, *J. Theor. Biol.* **266**, 11 (2010).
- [142] K. Harada and T. Ikegami, Evolution of specificity in an immune network, *J. Theor. Biol.* **203**, 439 (2000).
- [143] J. Hopfield, Physics, computation, and why biology looks so different, *J. Theor. Biol.* **171**, 53 (1994).
- [144] M. D. Kirchhoff and T. Froese, Where there is life there is mind: In support of a strong life-mind continuity thesis, *Entropy* **19**, 169 (2017).
- [145] S. Manicka and M. Levin, The cognitive lens: A primer on conceptual tools for analysing information processing in developmental and regenerative morphogenesis, *Philos. Trans. R. Soc. B* **374**, 20180369 (2019).
- [146] S. E. Marzen and J. P. Crutchfield, Optimized bacteria are environmental prediction engines, *Phys. Rev. E* **98**, 012408 (2018).
- [147] A. Mitchell, G. H. Romano, B. Groisman, A. Yona, E. Dekel, M. Kupiec, O. Dahan, and Y. Pilpel, Adaptive prediction of environmental changes by microorganisms, *Nature (London)* **460**, 220 (2009).
- [148] M. Sorek, N. Q. Balaban, and Y. Loewenstein, Stochasticity, bistability and the wisdom of crowds: A model for associative learning in genetic regulatory networks, *PLoS Comput. Biol.* **9**, e1003179 (2013).
- [149] G. Tkačik and W. Bialek, Information processing in living systems, *Annu. Rev. Condens. Matter Phys.* **7**, 89 (2016).
- [150] R. A. Watson, R. Mills, C. Buckley, K. Kouvaris, A. Jackson, S. T. Powers, C. Cox, S. Tudge, A. Davies, L. Kounios *et al.*, Evolutionary connectionism: Algorithmic principles underlying the evolution of biological organisation in evo-devo, evo-eco, and evolutionary transitions, *Evol. Biol.* **43**, 553 (2016).
- [151] S. J. Bartlett and P. Beckett, Probing complexity: Thermodynamics and computational mechanics approaches to origins studies, *Interface Focus* **9**, 20190058 (2019).
- [152] T. Froese, J. I. Campos, K. Fujishima, D. Kiga, and N. Virgo, Horizontal transfer of code fragments between protocells can explain the origins of the genetic code without vertical descent, *Sci. Rep.* **8**, 3532 (2018).
- [153] D. Michel, Life is a self-organizing machine driven by the informational cycle of Brillouin, *Orig. Life Evol. Biosph.* **43**, 137 (2013).
- [154] S. I. Walker and P. C. Davies, The algorithmic origins of life, *J. R. Soc. Interface* **10**, 20120869 (2013).
- [155] S. I. Walker, Top-down causation and the rise of information in the emergence of life, *Information* **5**, 424 (2014).
- [156] D. Hebb, *The Organization of Behavior* (Wiley, New York, NY, 1949).
- [157] J. J. Hopfield, Neural networks and physical systems with emergent collective computational abilities, *Proc. Natl. Acad. Sci. USA* **79**, 2554 (1982).

- [158] J. J. Hopfield, Neurons, dynamics and computation, *Phys. Today* **47**, 40 (1994).
- [159] S. McGregor, V. Vasas, P. Husbands, and C. Fernando, Evolution of associative learning in chemical networks, *PLoS Comput. Biol.* **8**, e1002739 (2012).
- [160] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevE.106.034401> for more details on the associative network.
- [161] B. Shenhav *et al.*, Graded artificial chemistry in restricted boundaries, in *Proceedings of the Conference on Artificial Life IX, Boston, MA* (MIT Press, Cambridge, MA, 2004), p. 501.
- [162] <https://www.youtube.com/watch?v=g1-SMI4Cl-A>
- [163] A. B. Boyd, D. Mandal, and J. P. Crutchfield, Leveraging environmental correlations: The thermodynamics of requisite variety, *J. Stat. Phys.* **167**, 1555 (2017).
- [164] R. A. Brittain, N. S. Jones, and T. E. Ouldrige, Biochemical Szilard engines for memory-limited inference, *New J. Phys.* **21**, 063022 (2019).
- [165] S. Ito and T. Sagawa, Maxwell's demon in biochemical signal transduction with feedback loop, *Nat. Commun.* **6**, 7498 (2015).
- [166] R. Booth, Y. Qiao, M. Li, and S. Mann, Spatial positioning and chemical coupling in coacervate-in-proteinosome protocells, *Angew. Chem.* **131**, 9218 (2019).
- [167] C. G. De Kruif, F. Weinbreck, and R. de Vries, Complex coacervation of proteins and anionic polysaccharides, *Curr. Opin. Colloid Interface Sci.* **9**, 340 (2004).
- [168] T. Z. Jia, C. Hentrich, and J. W. Szostak, Rapid RNA exchange in aqueous two-phase system and coacervate droplets, *Orig. Life Evol. Biosph.* **44**, 1 (2014).
- [169] T. D. Tang, C. R. C. Hak, A. J. Thompson, M. K. Kuimova, D. Williams, A. W. Perriman, and S. Mann, Fatty acid membrane assembly on coacervate microdroplets as a step toward a hybrid protocell model, *Nat. Chem.* **6**, 527 (2014).
- [170] D. S. Williams, A. J. Patil, and S. Mann, Spontaneous structuration in coacervate-based protocells by polyoxometalate-mediated membrane assembly, *Small* **10**, 1830 (2014).